Tina Memo No. 2010-007
Technical Memo

The TINA Geometric Morphometrics Toolkit

P.A. Bromiley, H. Ragheb and N.A. Thacker

Last updated
9 / 2 / 2012

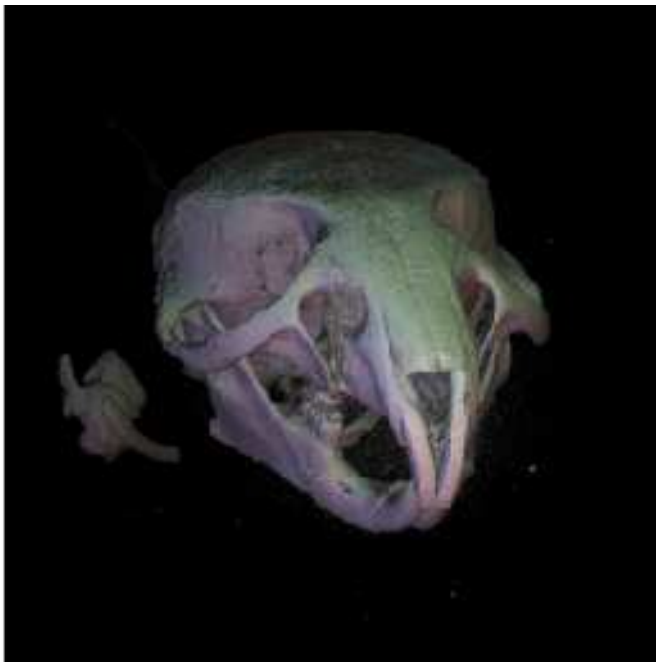

WWW.TINA-VISION.NET

Imaging Science and Biomedical Engineering Division,
Medical School, University of Manchester,
Stopford Building, Oxford Road,
Manchester, M13 9PT.

The TINA Geometric Morphometrics Toolkit

P.A. Bromiley, H. Ragheb and N.A. Thacker
Imaging Science and Biomedical Engineering Division
Medical School, University of Manchester
Manchester, M13 9PT, UK
paul.bromiley@manchester.ac.uk

1 Introduction

The TINA Geometric Morphometrics Toolkit comprises three main tools:

- the Manual Landmark tool;
- the Automatic Landmark Point Placement tool;
- the Shape Analysis tool.

The TINA Manual Landmark tool is designed to support the identification and specification of morphological landmarks for subsequent processing using techniques such as Procrustes analysis. However, the time required to manually identify sufficient numbers of landmarks in sufficient numbers of specimens for landmark-based geometric morphometric studies can be prohibitive. Therefore, the Automatic Landmark Point Placement tool can, given multiple examples of manually identified landmarks, semi-automatically identify those landmarks in further specimens, greatly reducing the user input required and therefore increasing productivity. The Shape Analysis tool provides standard Procrustes alignment of sets of manually or automatically identified landmark points, together with a statistically valid alternative to Procrustes alignment based on likelihood and taking full account of the anisotropic errors on individual landmarks.

This user manual covers:

- Section 2: installation of TINA, the Volpack volume rendering library, and the TINA Geometric Morphometrics toolkit on Linux systems.
- Section 3: the TINA graphical subsystem, TVs and TV tools.
- Section 4: the TINA Sequence tool, used for loading 3D medical data sets.
- Section 5: the Manual Landmark tool, definition of coordinate systems, and the TVs available in the tool.
- Section 6: the TINA interface to the Volpack volume rendering library, used to produce the volume rendered images displayed in the 3D TV of the Manual Landmark tool.
- Section 7: the Landmark Points dialog box, used to load, save and manipulate lists of landmark points.
- Section 8: the use of the Manual Landmark tool to mark up the positions of landmarks specified in the landmark point list.
- Section 9: the Keyboard Controls dialog box, used to assign keyboard shortcuts to the TVs of the Manual Landmark tool.
- Section 10: the TINA View tool, which provides functions to save the images displayed in TINA TV tools.
- Section 11: an example landmarking session.
- Sections 12 to 18: the Automatic Landmark Point Placement tool, which can semi-automatically identify landmarks given a database of manually identified examples.
- Sections 19 and 20: the Shape Analysis tool, used to analyse morphometric landmark points.
- Section 21: quick reference guides, explaining the functions of each button, data entry field etc. in the Sequence tool, Geometric Morphometrics toolkit, and Geometric Morphometrics toolkit dialog boxes.

This manual covers the Geometric Morphometrics Toolkit v. 2.0, which requires TINA 6.0 build number 008 or higher. Some familiarity with the Linux OS is assumed e.g. how to start and use a shell tool.

Throughout this guide, boxed sections contain detailed instructions on how to perform specific tasks with the TINA Geometric Morphometrics toolkit.

1.1 Citation

The authors would appreciate it if the following reference would be included in any published work that uses the TINA Geometric Morphometrics toolkit:

A.C. Schunke, P.A. Bromiley, D. Tautz and N.A. Thacker.
TINA manual landmarking tool: software for the precise digitization of 3D landmarks.
Frontiers in Zoology 9(1):6, 2012.

2 Installation

Installation of the TINA Geometric Morphometrics toolkit is performed in three stages: installation of the TINA libraries, which provide the TINA machine vision functionality; installation of the Volpack library, which is used by the TINA Manual Landmark tool to produce 3D volume rendered images of tomographic image data; and installation of the TINA Geometric Morphometrics toolkit, which provides the Manual Landmark, Automatic Landmark Point Placement, and Shape analysis tools and the user interface.

2.1 Install GTK+ 2

The TINA libraries depend on the GTK+ 2 libraries, which must be installed prior to installing TINA. Both the gtk2 and gtk2-devel packages are required: see the documentation for the Linux distribution you are using for instructions on how to install them. Installation of these packages should also result in installation of various other packages (atk, pango, fontconfig, pkgconfig) required by both GTK+ 2 and TINA.

2.2 Install the TINA libraries

Download the latest tina-libs and tina-tools tarballs from `http://www.tina-vision.net/tarballs/`. The TINA Geometric Morphometrics toolkit version 2.0 requires rcbuild008 or greater.

Note that some sections of the Geometric Morphometrics toolkit, particularly the volume renderer and the Automatic Landmark Point Placement tool, are extremely processor intensive. In order to mitigate the resultant time requirements, the software makes extensive use of parallel processing through OpenMP. However, this is not enabled by default, since it is not available in all compilers (any version of GCC from 4.2 onwards, released May 13th 2007, includes OpenMP); the user must specifically select it. In order to do so, specific compiler options must be used when the TINA libraries and the Geometric Morphometrics toolkit are built. If OpenMP is available in the compiler being used, enable parallelisation using the -fopenmp option by replacing all occurrences of

```
./configure
```

in the following instructions to

```
./configure CFLAGS="-O2 -fopenmp"
```

(the -O2 flag switches on the optimiser, which also speeds up the execution of the software). In addition, the main limitation on the execution speed of the automatic landmark point location software, when built with parallelisation options switched on, is not the number of available execution cores but the memory bandwidth. Therefore, the execution of the software will be slowed down if hyperthreading is used and it is recommended that, if hyperthreading is provided by the processor of the machine used to run the software, it should be disabled (this can be done through the bios setup programme, although the exact procedure will vary between computers).

To build the TINA libraries, open a shell tool and, as root, create a directory called Tina6 in /usr/local. Copy both tarballs into the new directory. Then unzip both tarballs using e.g.

```
gunzip tina-libs-6.0rcbuild008.tar.gz
gunzip tina-tools-6.0rcbuild008.tar.gz
```

and then unpack both using

```
tar -xvf tina-libs-6.0rcbuild008.tar
tar -xvf tina-tools-6.0rcbuild008.tar
```

Build the tina-libs libraries first: cd into the tina-libs-xxx (where xxx is the build number) directory, created when the tina-libs tarball was unpacked, and type

```
./configure
make
make install
```

Build the tina-tools libraries after the tina-libs libraries. cd into the tina-tools-xxx (where xxx is the build number) directory, created when the tina-tools tarball was unpacked, and type

```
./configure
make
make install
```

Note that the configure script in the tina-tools directory provides some information about the required libraries: in particular, if the necessary GTK+ 2 libraries have not been installed (see Section 2.1) or cannot be found, a warning will be issued.

2.3 Install Volpack

The TINA Geometric Morphometrics toolkit requires the Volpack volume rendering library. A version of this library can be found at `http://www.tina-vision.net/tarballs/manual_landmark_toolkit`. Unzip and unpack the tarball (this can be done in any directory, although it is standard practice to put user-installed software in /usr/local), cd into the new Volpack directory created when the tarball was unpacked, and use

```
./configure
make
make install (as root)
```

to build the library. On standard Linux systems the Volpack library and header file will be copied into standard locations (/usr/local/lib and /usr/local/include respectively). Then, when the TINA Geometric Morphometrics Toolkit is built, the compiler should automatically pick up the location of the header file, and the location of the library will be automatically picked up if /usr/local/lib is specified in /etc/ld.so.conf (the system-wide list of library locations). If it is not, then add /usr/local/lib to the file. Run ldconfig to refresh the list of libraries after the installation of Volpack (or restart the machine).

2.4 Install the Geometric Morphometrics Toolkit

Download the TINA Geometric Morphometrics toolkit (max_planck_toolkit-2.0.tar.gz) from the TINA web-site `http://www.tina-vision.net/tarballs/manual_landmark_toolkit`, place it in a convenient directory (e.g. `/home/my_username/tina_toolkits`), then unzip and unpack it. A new sub-directory will be created: cd into this subdirectory and type

```
./configure
```

The configure script will search for the external libraries needed to build the toolkit, issuing warnings about any it fails to find. In particular, if it cannot find the tina-libs or tina-tools libraries and headers, then their locations can be passed to the configure command using the following flags

```
–with-tina-includes=PATH
–with-tina-libraries=PATH
–with-tinatool-includes=PATH
–with-tinatool-libraries=PATH
```

e.g. ./configure –with-tina-libraries=/usr/local/Tina6/tina-tools/lib

If the tarballs were unpacked in /usr/local/Tina6, then the correct locations are e.g.

```
–with-tina-includes=/usr/local/Tina6/tina-libs-6.0rcbuild008
–with-tina-libraries=/usr/local/Tina6/tina-libs-6.0rcbuild008/lib
–with-tinatool-includes=/usr/local/Tina6/tina-tools-6.0rcbuild008
–with-tinatool-libraries=/usr/local/Tina6/tina-tools-6.0rcbuild008/lib
```

Finally, type “make” to build the toolkit.

The TINA Geometric Morphometrics toolkit needs to find the Volpack header and library. On standard Linux systems, these should be located in /usr/local/include and /usr/local/lib once Volpack has been installed. The header file should be picked up automatically by the compiler, and the library location should be automatically identified if /usr/local/lib is listed in /etc/ld.so.conf (and if ldconfig has been run, or the machine rebooted, since Volpack was installed). If these files cannot be located, the build will fail with error messages about functions with names starting with vp e.g. vpRenderRawVolume (these are the functions provided by the Volpack library). In that case, the locations of the Volpack library and header file can be specified manually to the configure command using the flags.

–with-volpack-include=PATH
–with-volpack-library=PATH

The executable file for the toolkit is located in the src subdirectory; alternatively, typing “make install” will copy it to the bin subdirectory. In either case, cd to that directory and type

./tinaTool

to run the toolkit.

Upon typing ./tinaTool, the top-level TINA tool will appear on the screen. This provides access to various sub-tools that in turn provide access to the algorithmic functionality. The sub-tools of most interest for geometric morphometrics are the Manual Landmark tool itself and the Automatic Landmark Point Placement and Shape Analysis tools, which can be accessed through the Manual Landmark tool, the sequence tool (used for loading medical image data) and the TV tools (used to display graphics). The text window at the bottom of the top-level TINA tool is used to display information and error messages.

3 TVs and TV Tools

Figure 1: A TV Tool.

The TINA image display system is designed to provide the maximum range of configuration options to the user. It consists of two components: TVs and TV tools. The majority of the TINA sub-tools will have a choice list at the top of the sub-tool window labelled "Tv:". Each button in the list represents an available graphics channel; however, TVs are not displayed by default, in order to save space on the screen.

In order to display a TV graphics channel, it must be associated with a TV tool:

- Start a new TV tool by pressing the "New Tvtool" button in the top-level TINA tool. A new TV tool window will appear on the screen, as shown in Fig. 1.
- Select the TV you wish to display, by left-clicking on its field in a "Tv:" list. Note that, since multiple lists are available in various TINA sub-tools, this may be required even if it appears that the desired field is already selected.
- Click the "install" button in the new TV tool. The title of the TV tool will change to the name of the TV it is displaying.

Once a TV has been installed on a TV tool, the TV tool will display the graphics produced by that TV, automatically updating whenever the graphics change without further user interaction. TV tools can be minimised or closed at any time without affecting algorithmic functionality. TV tools can also be freely reassigned, by selecting another TV and pressing the "install" button again. Note that, when a TV is installed onto a TV tool, any other TV tool displaying that TV will be un-installed: the title of the TV tool will change to reflect this.

Each TV tool has the following menus and buttons:

- Size: set the size of the TV tool to various default values (256x256, 512x512 and 768x768). TV tools can also be freely resized by clicking and dragging the corner of the TV tool window with the mouse.
- Mouse: sets the mouse interaction with the image in the TV tool.
- ROI: select a region of interest within the image.
- Proj: select the projection function for the TV tool.
- install: install the currently selected TV onto the TV tool.
- clone: install a duplicate of the currently selected TV onto the TV tool (allows multiple TV tools to display the same TV).
- init: re-initialise the image display to the default viewpoint, removing any translation, rotation or zooming of the image.

- repaint: re-display the image with the current translation, rotation or zooming (used if the image display should become corrupted for any reason).

See the TINA User's Guide [19] for full details.

Each TV tool also displays two messages: the name of the TV installed onto the TV tool, displayed in the title bar, and the current mouse functionality for the TV tool, displayed just below the button row. The mouse functionality lists the name of the function type followed by the functionality of the left, middle and right mouse buttons e.g. zoom3D:rot/zoom/shift indicates that the mouse has been set for 3D interaction with the image displayed in the TV tool, clicking and dragging with the left mouse button rotates the image in 3D, clicking and dragging with the middle mouse button zooms the image, and clicking and dragging with the right mouse button moves the image around in the TV tool.

4 Loading 3D Medical Image Data: The Sequence Tool

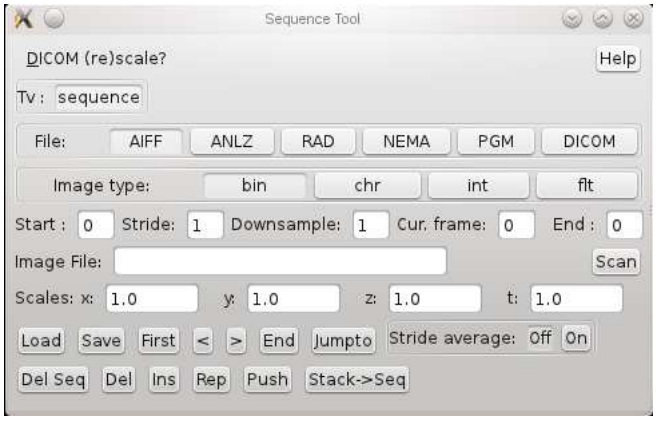
Sequence Tool
DICOM (re)scale?
Help
Tv : sequence
File: AIFF ANLZ RAD NEMA PGM DICOM
Image type: bin chr int flt
Start : 0 Stride: 1 Downsample: 1 Cur. frame: 0 End : 0
Image File:
Scan
Scales: x: 1.0 y: 1.0 z: 1.0 t: 1.0
Load Save First < > End Jumpto Stride average: Off On
Del Seq Del Ins Rep Push Stack->Seq

Figure 2: The Sequence tool, used for loading 3D medical image data.

The TINA Manual Landmark tool [18] allows manual identification of anatomical landmarks in 3D medical image data; it was designed for use with micro-CT images of rodent skulls, but can interact with any other tomographic medical image data e.g. MR images. Tomographic medical image data is loaded into TINA using the sequence tool, shown in Fig. 2.

The current sequence can be viewed by installing a TV tool onto the "Sequence" TV. This displays the slice specified in the "Curr. frame" field. When a sequence is loaded, the current slice is set to be the first slice in the sequence. Pressing the > button increments the current slice number i.e. moves forwards through the sequence. Pressing the < button decrements the current slice number i.e. moves backwards through the sequence. The "First" and "Last" buttons set the current slice to be the first and last slice of the sequence respectively. To view a specific slice, enter the slice number into the "Curr. frame" field and press the "Jumpto" button.

Individual image slices can be manipulated using the buttons at the bottom of the Sequence tool. The "Push" button copies the current slice to the top of the Imcalc stack (see the TINA User's Guide [19]) where it can be processed or passed to other tools. The "Ins" button takes the image from the top of the Imcalc stack and places it into the sequence after the current slice (if no sequence is loaded, then a new sequence will be created with the image as the first slice). The "Rep" button takes the image from the top of the Imcalc stack and replaces the current slice of the sequence with it. The "Del. seq" button deletes the entire current sequence, allowing it to be replaced with other images from the Imcalc stack. The "Stack->Seq" button will take all of the images currently held in the Imcalc stack and use them to create a new sequence (any current sequence is deleted in this process); the images will be placed into the sequence in reverse order i.e. the image on the top of the stack will become the last image in the new sequence.

In order to load a tomographic image data set:

- Enter the pathname of the image file/files into the "Image file:" field of the tool, either directly or using the file browser started by pressing the "Scan" button.
  - For AIFF, RAD, NEMA and PGM file types the sequence tool will automatically add the correct file name extension (.aiff, .rad, .ani and .pgm respectively), and so the filename should be entered without the extension. However, for Analyze or DICOM files the sequence tool will not automatically add an extension, and so the correct extension must be entered into the filename field. The file-browser started using the "Scan" button will automatically strip the extension from any selected file and so, when loading DICOM files, the extension must be manually re-entered into the filename field.
- Select the format of the image file from the "File:" choice list.
  - For file types (e.g. Analyze) where multiple slices of data are contained within a single file, enter the number of slices in the file into the "End" field. If you are unsure of the number of slices, then enter a number that is too high: the Sequence tool will load data until it reaches the end of the file and then reset the "End" field to show the number of slices loaded.
  - For file types (e.g. AIFF) where each slice is contained in a separate file, use the # character as a wildcard. For example, if the directory contains 10 files named im-00 to im-10, enter the pathname of the first file into the "Image file:" field, either by typing or by using the file-browser. Then replace the part of the name that changes between files with the # character, so that it becomes im-##. Enter the replaced part of the filename from the first file into the "Start" field, including any leading zeros (00 in this example) and the relevant part of the filename from the last file into the "End" field (10 in this case). During image loading, the wildcard characters in the pathname will be replaced with the contents of the "Start" field, and the number contained there will be incremented until it is equal to or greater than the "End" field, loading a file each time. If the numbers in the filenames increase by more than 1 between each file, then the "Stride" field can be used to specify the increment. For example, entering im-## into the "Image file:" field, 00 into the "Start" field, 10 into the "End" field and 2 into the "Stride" field will load the images im-00, im-02, im-04, im-06, im-08 and im-10 in that order. Note that there are some restrictions on loading multiple files in this way. The substitution only works for numbers, can only be applied to one contiguous segment of the pathname, and only works in ascending or descending order. In addition, all files must be in the same directory.
- If the sequence is too large to load into memory, the "Stride" and "Downsample" fields can be used to down-sample the data as it is loaded. The "Stride" field down-samples data in the inter-slice direction. For example, if the "Stride" field is set to 2, the sequence is down-sampled by a factor of 2 in the inter-slice direction. The "Stride average" switch controls the method used to achieve this; if it is set to "Off", then down-sampling is achieved by skipping over slices during loading. No use is made of the skipped slices; they are never loaded. If the "Stride average" switch is set to "On", then all slices are loaded and linear averaging is applied in the inter-slice direction to down-sample the sequence. The "Downsample" field down-samples data within each slice. For example, if the "Downsample" field is set to 2, then each image is down-sampled by a factor of 2 along its x and y axes after it is loaded, and the down-sampled image is entered into the sequence. Linear averaging is used during this process, reducing image noise.
- Press the "Load" button. Any previous image sequence will be deleted from memory, and the file(s) specified in the "Image file:" field will be loaded in order to produce a new sequence. The "Image type:" choice list will be updated to show the variable type of the images loaded into the sequence. The "Start" and "End" fields will be updated to show the slice numbers found during loading. The "Scales" fields will also be updated to show the size of each voxel, if this information is contained in the image file headers.

Note that, when the value entered into the "Stride" field is greater than 1 and the "Stride average" switch is on, the software will load slices in multiples of "Stride" prior to averaging them together. If the number of slices in the image volume is not an exact multiple of "Stride" then a warning will be issued during the final loading operation, indicating that there are not enough slices to perform the final averaging operation, and the last few slices will be discarded. The "End" field will be updated to indicate how many slices have been loaded.

The size of the original image volume, together with the options used during loading i.e. the values entered into the "Stride" and "Downsample" fields, and the choice of the "Stride average" method, will affect the memory requirements and execution speed of the Manual Landmark tool, and the accuracy of the Automatic Landmark Point Placement tool. In particular, larger volumes will slow down user interaction with the 3D TV of the Manual

Landmark tool (see Sections 5.2.2 and 6), as the volume renderer will require more processor time to render an image. The software was tested during development using micro-CT images of rodent skulls, with original resolutions of 658×658 voxels within each slice, and consisting of 1000 to 1200 slices per volume. These were down-sampled by a factor of 2 along all axes of the volumes during loading i.e. the "Stride" and "Downsample" fields were both set to 2. This produced volumes of 329×329 voxels in each of 500 to 600 slices after loading. These figures can be used as a guide when setting the values of the "Stride" and "Downsample" fields, as the volumes are small enough to be loaded onto any machine with 2Gb or more system memory, and to allow the volume renderer to run in interactive time on typical hardware at the time of writing i.e. Intel Core 2 Duo (or equivalent) or better.

Switching the "Stride average" option to "On" will increase the time required to load an image volume since, rather than loading a sub-set of the slices, all slices will be loaded and a linear averaging process applied to produce the down-sampled sequence. However, this also has the effect of reducing the image noise on the volume, producing higher quality renderings in the 3D TV (see Section 5.2.2). More importantly, the Automatic Landmark Point Placement tool will produce more accurate results, and the error analysis stage of that software will be more reliable. Therefore, it is recommended that the stride averaging be switched on when using this software if "Stride" is set to a value greater than 1.

The sequence tool can also save data to file, although not all file formats are supported (this is to prevent the user from saving sequences where essential header data is not available, or is no longer valid due to processing that has been applied to the images).

In order to save a tomographic image data set:

- Enter the pathname of the image file/files into the "Image file:" field of the tool,
- Select the format of the image file from the "File:" choice list (AIFF, RAD, ANLZ and PGM are supported: NEMA and DICOM are not supported).
  - For file types (e.g. AIFF) where each slice is contained in a separate file, use the # character as a wildcard. For example, if the sequence contains 10 slices, to be saved to files named im-00 to im-10, enter the pathname of the first file into the "Image file:" field. Then replace the part of the name that changes between files with the # character, so that it becomes im-##. During image saving, the wildcard characters in the pathname will be replaced with the slice numbers of the images in the sequence. If the numbers in the filenames increase by more than 1 between each file, then the "Stride" field can be used to specify the increment. For example, entering im-## into the "Image file:" field and 2 into the "Stride" field will save the images im-00, im-02, im-04, etc. Note that there are some restrictions on saving multiple files in this way. The substitution only works for numbers, can only be applied to one contiguous segment of the pathname, and only works in ascending or descending order.
- Press the "Save" button.

Part 1: The Manual Landmark Tool

P.A. Bromiley
Imaging Science and Biomedical Engineering Division,
Medical School, University of Manchester,
Stopford Building, Oxford Road,
Manchester, M13 9PT.

5 The Manual Landmark Tool

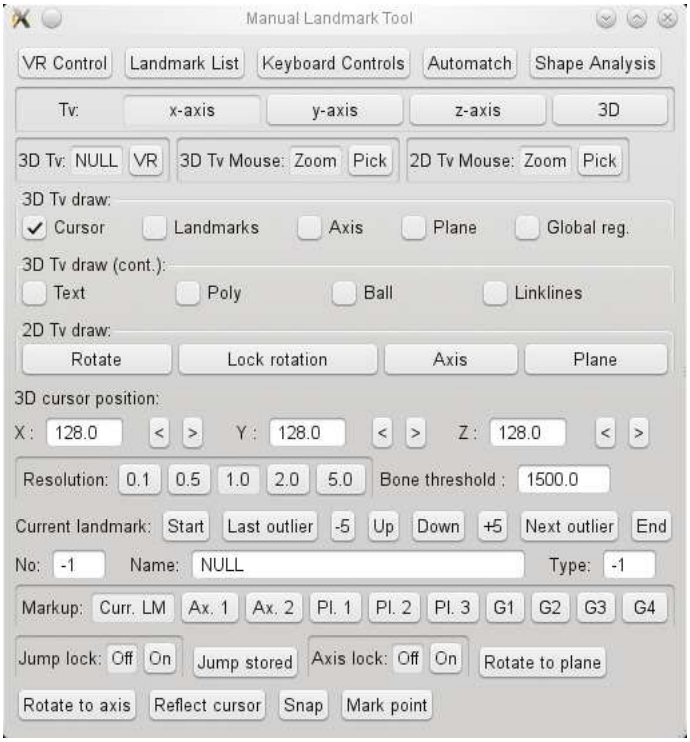

Figure 3: The Manual Landmark Tool.

The Manual Landmark tool (Fig. 3) [18] allows morphological landmarks to be identified within any 3D data set loaded into the Sequence tool. It provides multiple views of the data; three 2D TVs, showing orthogonal slices through the sequence, and a 3D TV showing a volume rendering of the whole sequence. The functionality of the Manual Landmark tool is based around the concept of a 3D cursor: a single point within the volume, whose coordinates can be manipulated by the user in a number of ways. The aim is to move the 3D cursor to the point the user wishes to mark, and then to store the coordinates of the cursor in a list of landmark points. When all landmark points have been specified in this way, the list of points can be saved to a file for further processing, e.g. Procrustes analysis, and/or added to the database used by the Automatic Landmark Point Placement tool.

5.1 Definition of Coordinates

The coordinate system used in the TINA Manual Landmark tool is derived from the original data, and relates to the images displayed by the sequence tool as follows: the z dimension is the inter-slice dimension, such that $z = z_{min}$ is the first slice and $z = z_{max}$ is the last slice. Within each image, the x dimension runs from left to right, such that $x = x_{min}$ is the left-hand edge of the image and $x = x_{max}$ is the right-hand edge, and the y dimension runs from top to bottom, such that $y = y_{min}$ is the top edge of the image and $y = y_{max}$ is the bottom. This definition is visible to the user in several places, such as in the definitions of the 2D TVs and in the landmark coordinates, and is shown in Fig. 4.

When loading large medical image volumes, it may be necessary to use the "Downsample" and "Stride" fields of the sequence tool, depending on available memory. In this case, it should be noted that the 3D cursor position will always be displayed in the coordinates of the down-sampled volume, but all landmark positions displayed in the Landmark Points dialog box (and saved to file using that dialog) will automatically be multiplied by the relevant down-sampling factors so that they are in the coordinate system of the original data. For example, if the original image volume consists of 1000 slices, each 1000x1000 voxels in size, and this is down-sampled by a factor of 2 on each axis when it is loaded (i.e. the "Downsample" and "Stride" fields in the sequence tool are both set to 2) then the coordinates of the 3D cursor placed at the centre point of the volume will be (256, 256, 256), and this will be displayed in the X,Y and Z fields of the Manual Landmark tool; however, if this point is marked up as a landmark point, then the coordinates will automatically be multiplied by the down-sampling factors to become (512, 512, 512) before they are displayed in the Landmark Points dialog box or saved to file. This prevents the choice of

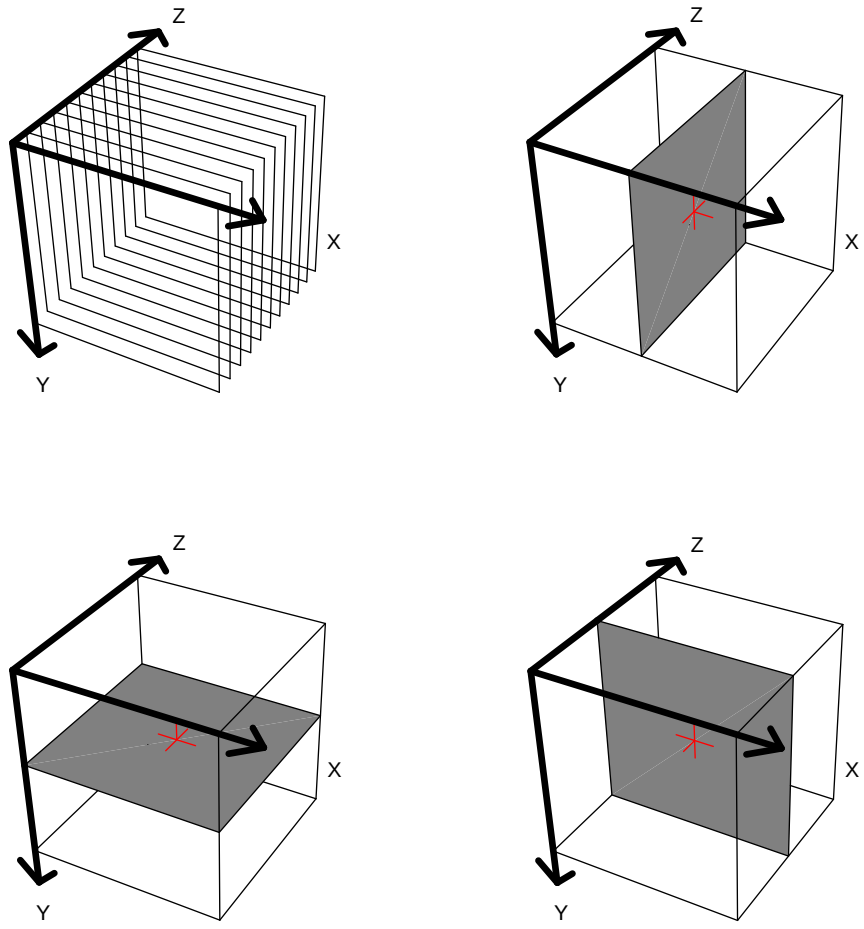

Figure 4: The coordinate system used to refer to the data (top left): the z-direction is the inter-slice direction; the x-direction is the left-right axis of images seen in the sequence tool TV, and the y-direction is the top-bottom direction of images seen in the Sequence tool TV. The images shown in the three 2D TVs of the Manual Landmark tool are generated from orthogonal planes passing through the current 3D cursor position (shown in red), such that the image shown in the x-axis TV is generated from a plane perpendicular to the x-axis (top right), the image shown in the y-axis TV is generated from a plane perpendicular to the y-axis (bottom left), and the image shown in the z-axis TV is generated from a plane perpendicular to the z-axis (bottom right).

down-sampling factors from affecting any subsequent Procrustes analysis of the landmark points.

5.2 Image Display

The Manual Landmark tool provides four TVs (see Section 3). Three of these provide 2D views of orthogonal slices through the volume, centred on the 3D cursor. The fourth provides a 3D volume rendering of the volume. The TVs can also display 3D cursor, landmark, axis and plane points. Figure 5 shows the four TVs in a recommended arrangement, such that the axes are consistent across the three 2D TVs.

5.2.1 2D Image Display

The three 2D TVs display orthogonal slices through the sequence, centred on the current 3D cursor position, which is displayed in each TV as a large, red cross-hair. Whenever the 3D cursor is moved, the images displayed in these TVs will be updated to reflect this change. The TVs are named according to the axis that is perpendicular to the image displayed in the TV (see Section 5.1), as shown in Fig. 4. Therefore, the x-axis TV displays an image in

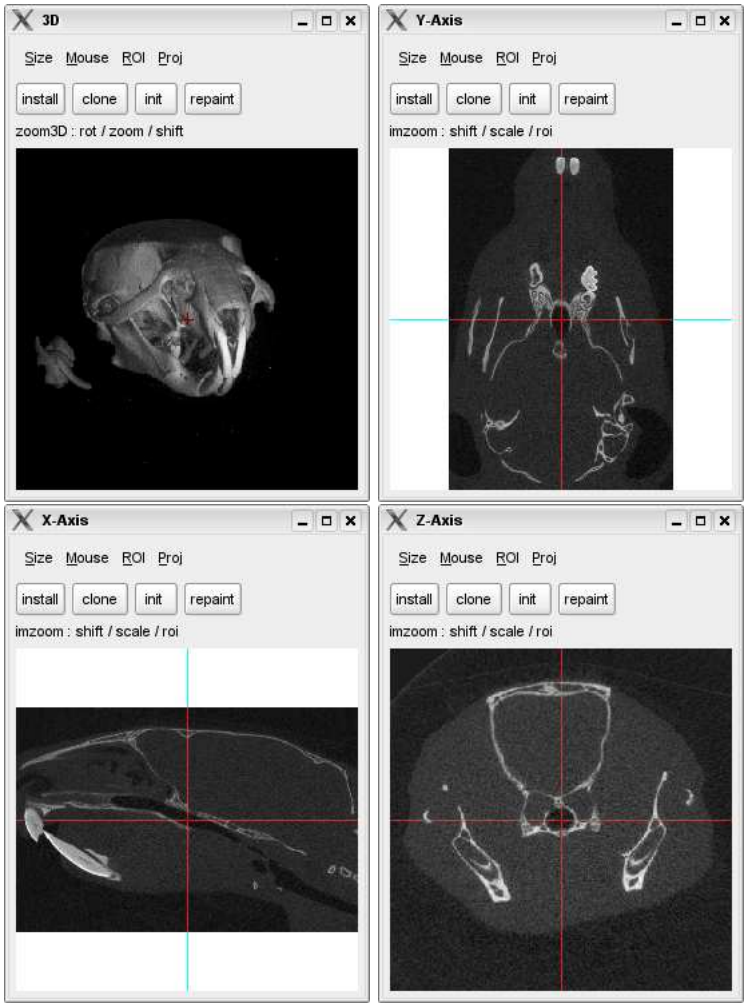
3D
Size Mouse ROI Proj
install clone init repaint
zoom3D : rot / zoom / shift
Y-Axis
Size Mouse ROI Proj
install clone init repaint
imzoom : shift / scale / roi
X-Axis
Size Mouse ROI Proj
install clone init repaint
imzoom : shift / scale / roi
Z-Axis
Size Mouse ROI Proj
install clone init repaint
imzoom : shift / scale / roi

Figure 5: The four TVs of the Manual Landmark tool, arranged so that the cross-hairs for the three 2D TVs are consistent across all three windows.

which the left-right axis is the z axis of the sequence and the top-bottom axis is the y axis of the sequence. The y-axis TV displays an image in which the left-right axis is the x-axis of the sequence and the top-bottom axis is the z-axis of the sequence. The z-axis TV displays an image in which the left-right axis is the x-axis of the sequence and the top-bottom axis is the y-axis of the sequence.

Mouse interaction with the 2D TVs can be selected using the "2D Tv Mouse:" switch in the Manual Landmark tool. The functions assigned to each mouse button are displayed in the TV tool, just above the image. When "Zoom" is selected, the mouse can be used to manipulate the image. Clicking and dragging with the left mouse button will move the image around in the TV. Clicking and dragging with the middle mouse button will zoom in or out. Clicking and dragging with the right mouse button will select a rectangular region of interest, and then zoom into that region when the mouse button is released. During the moving or zooming operations, the image will be replaced with a wire-frame in order to increase speed. When "Pick" is selected, the mouse can be used to manipulate the 3D cursor via the 2D TVs. Only the left mouse button has functionality assigned to it in this mode. Clicking on any point in the image will move the 3D cursor to that point: clicking and dragging will move the 3D cursor interactively. Only the coordinates displayed in the TV will be altered e.g. the x-axis TV can be used to manipulate the y and z coordinates of the 3D cursor.

When marking up large numbers of points in multiple images, it is desirable to keep the mouse pointer within the TVs (i.e. not to have to move it repeatedly back to the Manual Landmark tool window in order to click buttons there). Therefore, the TVs also provide a set of keyboard shortcuts for commonly used functions:

- Up arrow key: move the 3D cursor position upwards by one increment.
- Down arrow key: move the 3D cursor position downwards by one increment.

- Left arrow key: move the 3D cursor position left by one increment.
- Right arrow key: move the 3D cursor position right by one increment.
- Right ctrl: select "Zoom" mode for mouse interactions in all TVs.
- Right shift: select "Pick" mode for mouse interactions in all TVs.
- Enter: equivalent to clicking on the "Mark point" button in the Manual Landmark tool i.e. store the current 3D cursor coordinates as the currently selected landmark.
- Page Up: scroll up to the previous landmark in the landmark list (see Section 7).
- Page Down: scroll down to the next landmark in the landmark list (see Section 7).

The size of the increment for cursor movement is selected using the "Resolution" choice list in the Manual Landmark tool: by default it is set to one voxel.

Note that, in order for the TVs to receive the keyboard inputs, the window manager focus must be on the TV. In many window managers, the window that has the current focus is displayed with a coloured title bar: all other windows have a grey title bar. Most window managers provide two modes for focus selection: click-to-focus, in which the user must left-click with the mouse inside a window in order to switch focus to that window, and focus-follows-pointer, in which the focus in on whichever window the mouse pointer is inside. If click-to-focus is selected, then the user will have to left-click on a specific TV tool in order to send subsequent keyboard instructions to it. Therefore, if significant use is made of keyboard shortcuts, it may be easier to change the window manager focus policy to focus-follows-pointer, so that keyboard instructions are sent to whichever TV the mouse pointer is inside.

In order to switch to focus-follows-pointer mode in Suse 10.3:

- Click on the start button (in the bottom-left corner of the screen); select "Configure desktop".
- Select the "Desktop" tab from the menu on the left.
- Select the "Window Behaviour" tab from the menu on the left.
- Select the "Policy" drop-down menu in the "Focus" box.
- Select "Focus follows mouse" from the drop-down menu.
- Click the "Apply" button at the bottom of the window.

In order to switch to focus-follows-pointer mode in Suse 11.3:

- Click on the start button (in the bottom-left corner of the screen); select "System Settings".
- Select the "General" tab at the top of the System Settings dialog.
- Select the "Window Behaviour" tab from the "Look and Feel" section of the dialog.
- Select the "Window Behaviour" tab from the menu on the left-hand side of the dialog.
- Select "Focus follows mouse" from the "Policy" drop-down menu.
- Click the "Apply" button at the bottom of the window.

The window manager focus will now follow the mouse pointer. These instructions are valid for openSuse using the KDE desktop. Other window managers may vary: see the instructions for the window manager you are using if necessary.

The keyboard shortcuts can be reassigned to any set of keys (with the exception of a few reserved by the window manager e.g. the print screen key) using the TV Keyboard Controls dialog box (see Section 9).

Additional display functionality can be accessed via the "2D Draw" check list in the Manual Landmark tool. If "Axis" is selected, then any currently specified axis points (see Section 8.1) will be displayed as small yellow cross-hairs in the 2D TVs. Similarly, "Plane" will display any currently specified plane points (see Section 8.1) as small cyan cross-hairs. Note that these are projected onto the TVs, so that they are displayed even if they do not lie in the plane of the image.

The "2D Draw" check list in the Manual Landmark tool also provides two options linked to the display in the 3D TV (see 5.2.2). If "Rotate" is selected, then any rotation that the user applies to the image displayed in the 3D TV will also be applied to the images displayed in the 2D TVs. This allows the user to align the 3D image in a specific way (for example, looking along the anterior-posterior axis of the specimen with the plane of bilateral symmetry aligned with the vertical axis of the image) and then apply this rotation to the images shown in the 3D TVs to provide views of the data that are easier to interpret. However, any subsequent rotation of the image in the 3D TV will also be applied to the images in the 2D TVs. The "Lock rotation" option modifies this behaviour so that any rotation applied to the image displayed in the 3D TV, at the time the lock rotation option is selected, is stored and used for all subsequent drawing operations on the 2D TVs until the lock rotation option is deselected. The image in the 3D TV can then be manipulated without disrupting any useful alignment that has been achieved in the 2D TVs. The "Lock Rotation" option over-rides the "Rotate" option (i.e. changing the selection of the "Rotate" option has no effect whilst "Lock Rotation" is selected). When "Lock Rotation" is deselected, the images in the 2D TVs will be re-drawn either using the rotation currently applied to the image in the 3D TV, if "Rotate" is selected, or with no rotation (i.e. aligned to the axes of the original image volume) if "Rotate" is not selected.

When the "Rotate" or "Lock rotation" options are selected, the increment and decrement buttons (i.e. the arrow buttons) in the Manual Landmark Tool window still move the position of the 3D cursor along the major axes of the original image volume. However, when the keyboard shortcuts are used in the 2D TVs, the cursor is moved along the (rotated) axes of the image displayed in the TV, not along the axes of the original volume. This allows the user to rotate the volume to a preferred orientation and then select extremal points on structures displayed in the 2D TVs relative to that rotation. The user should be aware that the definitions of landmark points identified in this way is relative to the orientation they select. Therefore, although the orientation can be selected manually by rotating the image displayed in the 3D TV, it may be preferable to use a more rigorous definition. Section 8.1 describes methods to achieve this by identifying an arbitrary plane and axis using landmark points, and then rotating the image displayed in the 3D TV to a fixed orientation using these geometrical constructs.

Note that the "Rotate" and "Lock rotation" options modify the description of the coordinate system given in Section 5.1; the images displayed in the 2D TVs are no longer produced from planes orthogonal to the major axes of the image volume when either of these options are selected.

5.2.2 3D Image Display

The 3D TV displays a volume rendering of the data loaded into the Sequence tool. Since the preparation of this image involves a significant amount of processor time and memory usage, the volume renderer is not activated by default and must be explicitly switched on by the user. This is done using the "3D Tv" choice list in the Manual Landmark tool: if "NULL" is selected then the volume renderer is switched off. Selecting "VR" will switch on the volume renderer and prepare the sequence for rendering: this may take several tens of seconds to complete.

In order to use the volume renderer:

- Select "VR" in the "3D Tv" choice list in the Manual Landmark tool.
- Select "3D" in the "Tv:" choice list of the Manual Landmark tool, start a new TV tool, and click "Install" in the TV tool; the title bar of the TV tool will change to indicate that it is now displaying the 3D image stream.
- Click the "VR Control" button in the Manual Landmark tool to start the volume rendering controls dialog box.
- Use the volume rendering controls (see below) to produce a satisfactory rendering.

Mouse interaction with the 3D TV can be selected using the "3D Tv Mouse:" switch in the Manual Landmark tool. In "Zoom" mode, mouse interaction controls the 3D image. Clicking and dragging with the left mouse button will rotate the image. Clicking and dragging with the middle mouse button will zoom in or out (moving the mouse left and right) or rotate the image within the image plane (moving the mouse up or down). Clicking and dragging with the right mouse button will move the image around within the TV. When "Pick" mode is selected, mouse interaction within the 3D TV controls the 3D cursor. Left-clicking within the 3D TV will move the 3D cursor to the first bone surface under the mouse pointer i.e. the position of the mouse pointer specifies a vector through the data (depending on the current rotation of the 3D image); the software will run along this vector, starting at the uppermost point from the current view direction, until it reaches a bone surface. The 3D cursor will then be moved to this point on the bone surface. The bone surface is defined by a threshold, which the user can set in the "Bone threshold" field of the Manual Landmark tool (see Section 18). The software also provides a profiling tool

to aid the user in selecting a suitable threshold. Right clicking with the mouse within the 3D TV will produce a grey-level profile along the vector under the mouse pointer, and display it in the Imcalc Graph TV.

In order to select a bone threshold intensity:

- Start the volume renderer as described above. Using the "Zoom" mouse mode, rotate the image such that a point with a solid bone surface is clearly visible.
- Start the "Imcalc" tool from the top-level tinaTool.
- Click "Graph" in the TV choice list of the Imcalc tool: start a new TV tool, and click "Install" in the TV tool to assign it to the Imcalc Graph TV.
- Select "Pick" mode for the 3D TV mouse.
- Move the mouse pointer to a location within the 3D TV that has a clearly visible bone surface.
- Right click with the mouse: a grey-level profile through the image volume, along the vector under the mouse pointer, will be displayed in the Imcalc Graph TV. The x-axis of this graph represents depth through the 3D image, with 0 being the back of the image and 1 being the front of the image from the current viewpoint. The y-axis represents the intensity (grey level) of the data along the profile.
- Identify the step change in the intensity profile that represents the transition from soft tissue to bone. Several bone surfaces may be visible in the profile; also, be sure to clearly identify soft tissue from background.
- Calculate an intensity value half-way between the mean intensity of bone and the mean intensity of soft tissue.
- Enter the calculated intensity value into the "Bone threshold" field of the Manual Landmark tool.

As with the 2D TVs (see Section 5.2.1; the same comments about window manager focus apply to the 3D TV), several keyboard shortcuts are available in the 3D TV:

- Right ctrl: select "Zoom" mode for mouse interactions in all TVs.
- Right shift: select "Pick" mode for mouse interactions in all TVs.
- Enter: equivalent to clicking on the "Mark point" button in the Manual Landmark tool i.e. store the current 3D cursor coordinates as the currently selected landmark.
- Page Up: scroll up to the previous landmark in the landmark list (see Section 7).
- Page Down: scroll down to the next landmark in the landmark list (see Section 7).

These can be reassigned to any desired keys using the TV Keyboard Controls dialog box (see Section 9).

As well as displaying a volume rendered image of the data currently loaded into the sequence tool, the 3D TV can display any stored landmark, axis, plane or G-reg points (see Sections 7, 8.1 and 13.2.1). The user can choose which points to display using the "3D Tv draw" check list in the Manual Landmark tool. The available choices are:

- Cursor: display a red cross-hair representing the 3D cursor (switched on by default).
- Landmarks: display green cross-hairs for all marked-up landmark points. Any landmark point outside the boundaries of the currently loaded image volume will be displayed in purple instead of green (this is most likely to occur if landmarks corresponding to a different volume have been loaded from a file). Any automatically located landmark points that have failed an error check (see Section 12) will be displayed in green and red chequers.
- Axis: display yellow cross-hairs for any specified axis points.
- Plane: display cyan cross-hairs for any specified plane points.
- Global reg: display orange cross-hairs for any specified G-reg points.

- Text: display text labels for any displayed landmark, plane, axis or G-reg points. The text label will appear next to the point and will be displayed in the same colour (green for landmark points, yellow for axis points, cyan for plane points and orange for G-reg points). Landmark points are labelled with their point numbers, axis points as A1 and A2, plane points as P1, P2, and P3, and G-reg points as G1, G2, G3 and G4. Since a single point might be used simultaneously as a landmark, axis, plane and G-reg point, the labels are arranged around the point such that they should not overlap. However, only the labels corresponding to the selected drawing options will be displayed e.g. if a single point is identified both as a landmark point and an axis point, but the "Landmarks" option is selected and the "Axis" option is not selected, then the landmark number label will be displayed but the axis point label will not.
- Poly: display a yellow line representing the axis and a cyan grid representing the plane.
- Ball: display 3D spheres instead of cross-hairs for the landmark, axis, plane and G-reg points.
- Linklines: display green lines linking equivalent landmarks on the left and right sides of the sample, if the corresponding linking lines file has been prepared and loaded or the linking information has been manually entered into the Landmark Points dialog box (see Section 7).

Note that all points, whether displayed as cross-hairs or spheres, will be shaded to indicate their interaction with the bone surface as specified in the "Bone threshold:" field of the Manual Landmark tool. Any part of a cross-hair or sphere that lies above all bone surfaces from the current view direction will be displayed at 100% intensity. Any part of a cross-hair or sphere that lies below a bone surface from the current view direction will be displayed at 50% intensity. This provides visual information on how the point interacts with the surface.

6 The Volume Rendering Interface

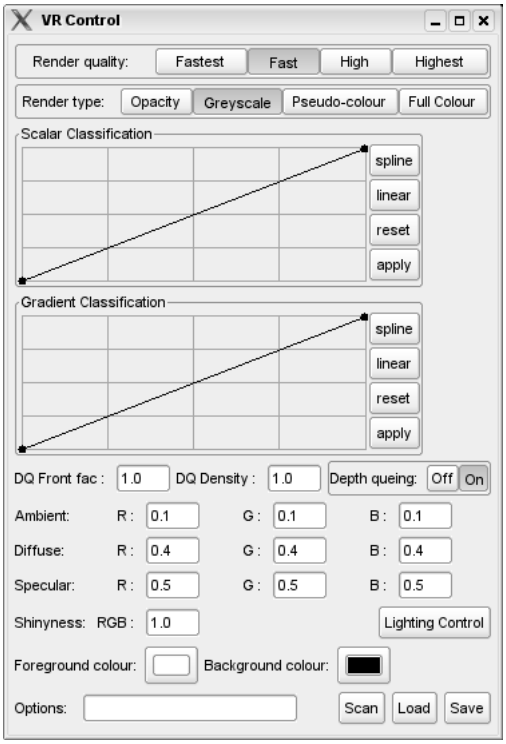

Figure 6: The interface to the volume rendering engine.

The TINA Manual Landmark tool uses the Volpack volume rendering library [11, 12] to produce 3D views of the data loaded into the Sequence tool. In general terms, it is easier to see the overall position of a morphological landmark when the data is displayed in 3D, although for maximum accuracy it is important to refine the landmark position by checking the 2D views, since these provide a more direct view of the data.

Pressing the "VR Control" button in the Manual Landmark tool spawns a dialog box that provides access to various user-controlled parameters of the volume rendering engine, as shown in Fig 6. The most important of these are the graphs contained in the "Scalar Classification" and "Gradient Classification" fields. Each of these takes the form of a small, user-controlled graph with four associated buttons. The curve shown in the "Scalar Classification" graph controls the way in which intensity (grey level) in the original 3D image data loaded into the Sequence tool (and displayed on the x-axis of the graph) is converted into the opacity of the corresponding pixel in the volume rendered image (the y-axis of the graph). The shape of the curve is controlled by a set of points placed within the graph by left-clicking with the mouse and moved by left-clicking and dragging with the mouse cursor placed over an existing point. The "spline" and "linear" buttons set the shape of the curve to be a spline or piecewise-linear curve through the control points. The "reset" button" returns the shape of the curve to the default linear ramp. After setting the curve to the desired shape, the "apply" button must be pressed to pass the new curve to the volume renderer. The effects of these parameters can be seen in Figs. 7, 8, 9, 10 and 11. In general, the curves should both be set to sigmoid functions: using steep sigmoids, as shown in Fig. 9, will produce a surface-rendering style image.

The balance between rendering speed and image quality is controlled by the "Render quality" choice list. These choices control an upper and lower threshold; voxels with intensities below the lower threshold are ignored during rendering, and voxels with intensities higher than the upper threshold are assigned 100% opacity (so that no voxels behind them need to be rendered). The "Fastest" choice sets these thresholds to their most extreme values, so that many voxels are ignored; the rendering is then faster at the expense of lower image quality. The "Highest" choice disables the thresholds completely, so that all voxels are rendered; the rendering is then slow but produces the best image quality. In general the "Fast" choice is the most suitable during mouse interaction with the image displayed in the 3D TV.

The volume renderer is also capable of producing several different types of image, depending on the selection in the "Render type" choice list. The "Opacity" choice will render an image showing only the opacity of each voxel i.e. with no surface shading, so that the "Gradient classification", lighting and material parameters are all ignored.

This provides a transparent view of the data as shown in Fig. 8. The "Greyscale" choice will render an image with both opacity and surface shading, so that the "Gradient Classification", lighting and shading parameters are all applied. However, only the overall intensity of light from each voxel is displayed and so the resulting image is greyscale. The "Pseudo-colour" choice produces a greyscale image but then colours the result using the colour chosen with the "Foreground colour" button, and sets the background to a uniform colour as specified with the "Background colour" button. The "Full colour" choice gives the renderer complete control of image colour, so that coloured lights and material properties can be used (see below); the "Background colour" field can also be used to change the colour of the background in this mode.

Pressing the "Lighting Control" button spawns the Lighting Control dialog box, which provides the user with full control over the lighting parameters used by the volume renderer. See Section 6.1 for more details.

The volume renderer also supports depth cueing. This function allows a user-controlled "fog" to be applied to the volume rendering, which grows thicker with depth through the image. More distant objects therefore appear dimmer, which helps to create a more three-dimensional effect. Depth cueing can be switched off or on using the "Depth cueing:" buttons, and is controlled by the two parameters "Front fac" (the thickness of the fog at the front of the image) and "Density" (the rate at which the fog grows thicker with depth through the image), according to

$$fog\ thickness = (front\ factor).e^{-density.depth}$$

Both "Front fac" and "Density" must be positive numbers. Changes to the these parameters are not implemented in the rendered image until the next time it is regenerated: the user can force the image to be rendered with the new parameters by clicking the "Repaint" button in the TV tool displaying the 3D TV.

When the rendering type is set to greyscale, pseudo-colour or full colour, the intensity gradient of the original 3D image data is used to shade the surface of each voxel. Finer control over the reflectivity of the surfaces can be achieved using the material property parameters contained in the "Ambient", "Diffuse", "Specular" and "Shininess" fields. The ambient and diffuse fields control how the surface reflects ambient and diffuse light respectively; the red, green and blue components can be set independently. The specular field controls the intensity of specular reflections (i.e. highlights), and again the red, green and blue components can be controlled independently. The shininess field controls the overall shininess of surfaces, for which there is only a single parameter. Valid parameters for the ambient, diffuse and specular RGB components are in the range 0.0 to 1.0: the shininess parameter should be in the range 1.0 to 100.0. The colouration of surfaces will only be visible when the renderer is in full colour mode; in greyscale or pseudo-colour modes the rendered image is produced using only the overall intensity of the reflected light (i.e. the intensity of each of the red, green and blue components combined). Although the Volpack library provides the ability to assign different material properties to different tissues, using the techniques described in [9], the current version of the Manual Landmarking tool assumes a single material type. Changes to the material properties are not implemented in the rendered image until the next time it is regenerated: the user can force the image to be rendered with the new parameters by clicking the "Repaint" button in the TV tool displaying the 3D TV.

It is important to note that none of the options in the VR Control or Lighting Control dialog boxes will affect the accuracy of landmark point placement (except to the extent that they affect the user's ability to see where the landmark point should be placed). When the user selects a point in the 3D image by placing the mouse cursor over that point, and then left-clicking with the "Pick" mode of the "3D Tv Mouse" function selected, the 3D cursor is moved to the first bone surface lying under the point specified by the mouse cursor. However, this is not done using the 3D renderer: instead, it is done by searching through the original 3D image data, as contained in the Sequence tool, until a voxel with an intensity equal to or greater than the intensity specified in the "Bone threshold" field of the Manual Landmark tool is found. None of the options contained in the VR Control dialog box are used in this process, and so the user is free to reset them in order to produce the best possible rendering of the data without this process having any effect on subsequent analysis of the landmark points generated using the software.

Significant user experimentation with the options displayed in the VR Control dialog box may be required in order to produce the best possible 3D display of the data. Therefore, the software provides the ability to save/load all of the currently selected volume rendering and lighting options to/from a file. The pathname of the file can be entered into the "Options:" field at the bottom of the VR Control dialog box, either by typing the pathname directly or by using the file browser started with the "Scan" button. The "Load" and "Save" buttons then load or save the options to or from the file. Several example volume renderer options files are available in the extra_files sub-directory of the max-planck-toolkit directory that is created when the Geometric Morphometrics Toolkit is extracted.

- tmlt_surface_style.txt - the options used in Fig. 9; these provide a surface-rendering style 3D image for many micro-CT scans of rodent skulls.

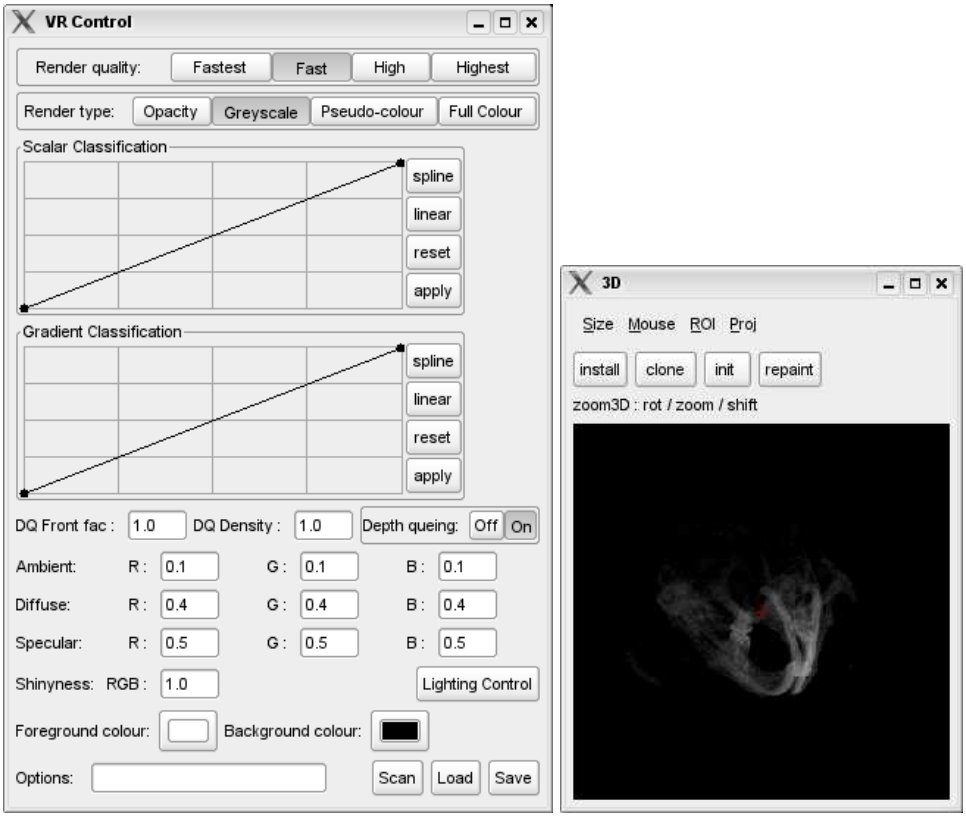

Figure 7: The volume renderer using the default options.

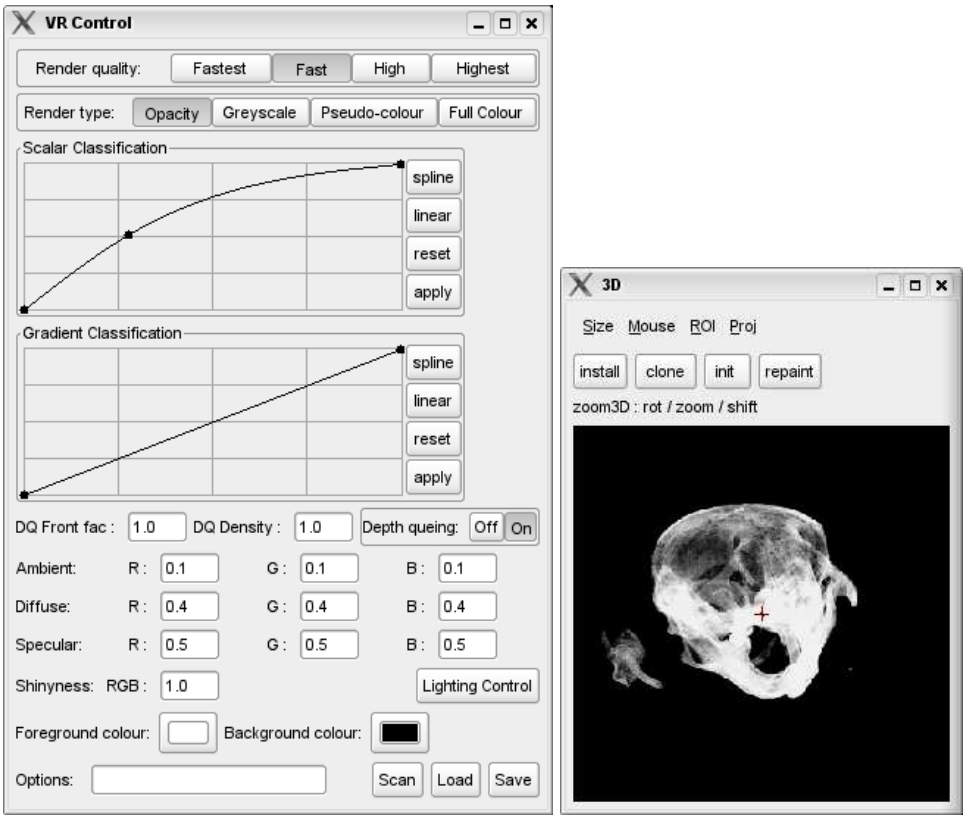

Figure 8: Opacity-only volume rendering.

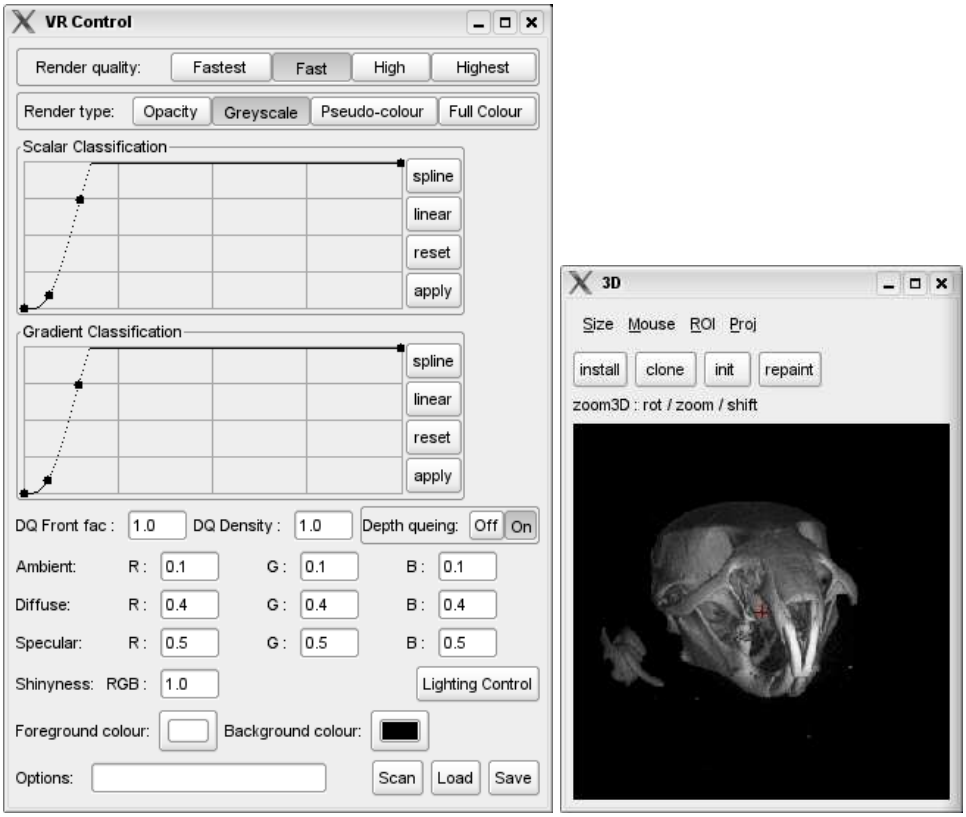

Figure 9: Greyscale volume rendering with options chosen to give a surface rendering style image.

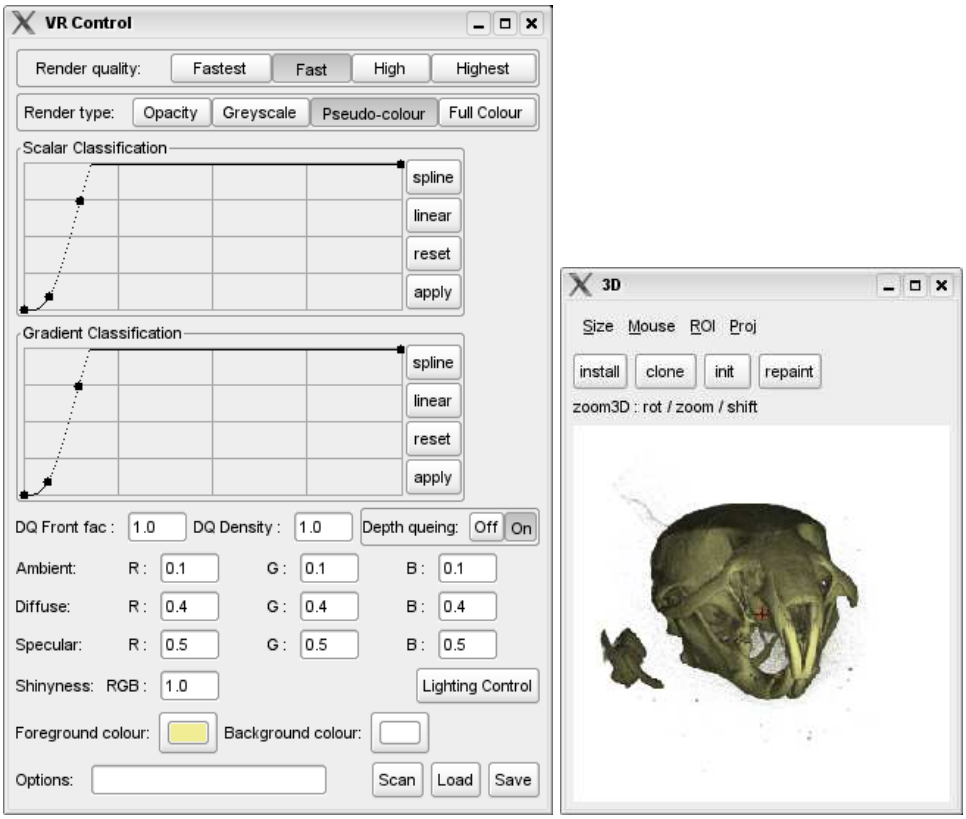

Figure 10: Pseudo-colour volume rendering with options chosen to give a surface rendering style image.

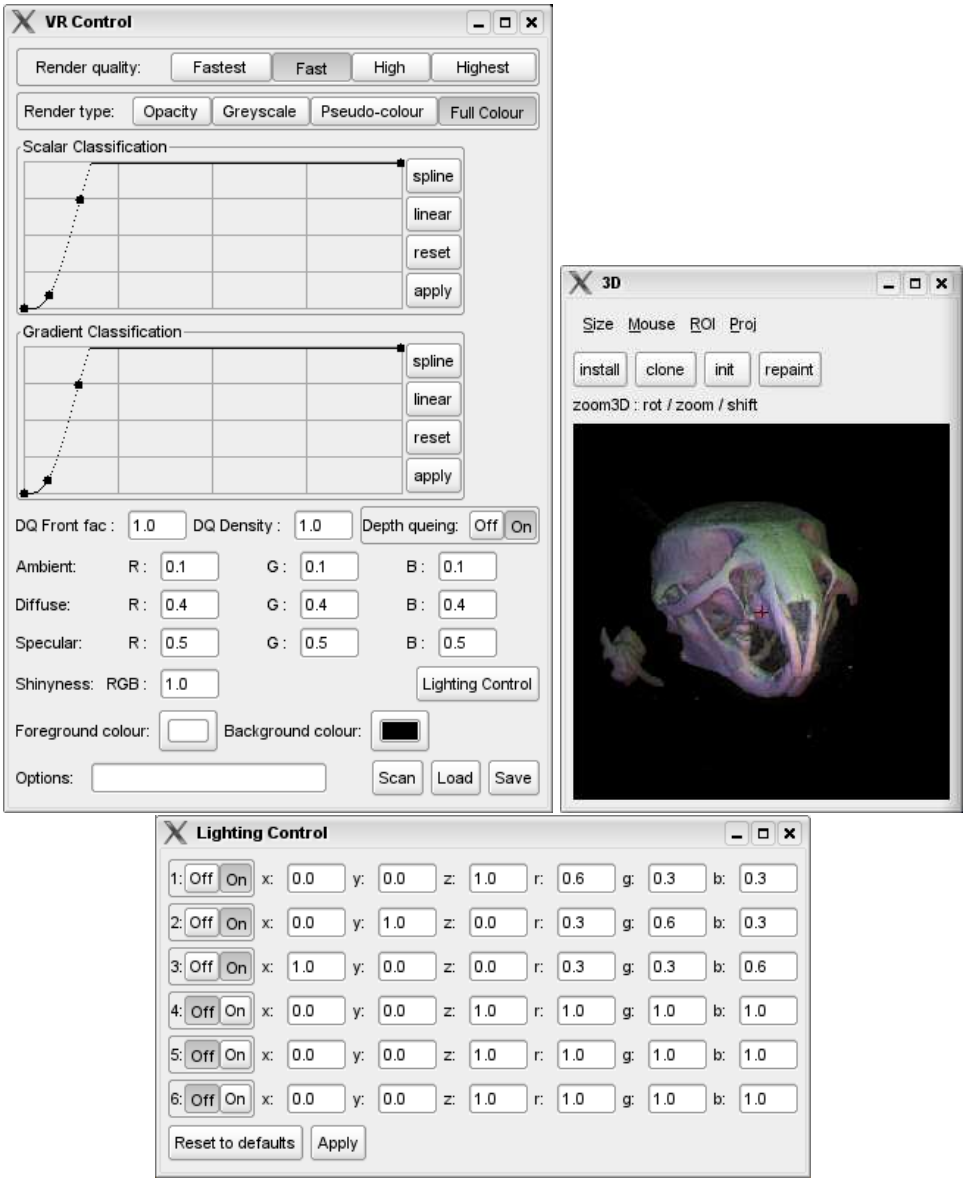

Figure 11: Full-colour volume rendering with options chosen to give a surface rendering style image, showing the effects of coloured lighting.

- tmlt_pseudo_colour.txt - similar to the options used in Fig. 9, but with pseudo-colour rendering enabled and a light-blue background, allowing the user to see the opacity of the rendered image more easily.
- tmlt_full_colour.txt - similar to the options used in Fig. 9, but with multiple coloured lights enabled; similar to the options used to produce the image on the front page of this document.
- tmlt_printing.txt - similar to the options used in Fig. 9, but with pseudo-colour rendering enabled with a white background, and with all landmark points displayed as spheres rather than cross-hairs: these options are useful for producing images to be incorporated into publications.

6.1 The Lighting Control Interface

The volume rendering engine supports up to six fully configurable light sources. Fig. 12 shows the lighting control dialog box, which allows the parameters of the lights to be specified by the user. The dialog box is created by pressing the "Lighting Control" button in the "VR Control" dialog box.

Each line of the lighting control dialog box controls one of the lights. The buttons on the far left specify whether the light is switched on or off. The next three fields (x, y and z) specify the normalised x, y, and z components of

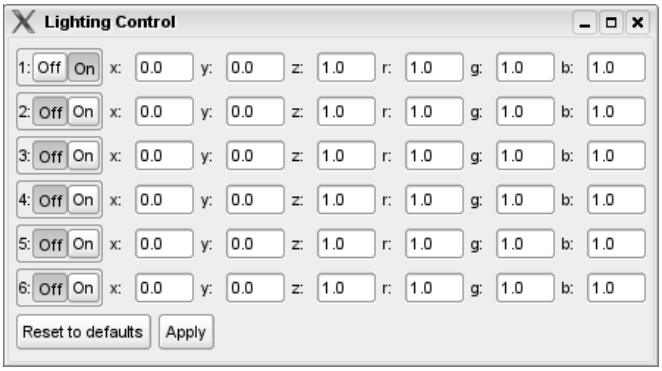
Lighting Control
1: Off On x: 0.0 y: 0.0 z: 1.0 r: 1.0 g: 1.0 b: 1.0
2: Off On x: 0.0 y: 0.0 z: 1.0 r: 1.0 g: 1.0 b: 1.0
3: Off On x: 0.0 y: 0.0 z: 1.0 r: 1.0 g: 1.0 b: 1.0
4: Off On x: 0.0 y: 0.0 z: 1.0 r: 1.0 g: 1.0 b: 1.0
5: Off On x: 0.0 y: 0.0 z: 1.0 r: 1.0 g: 1.0 b: 1.0
6: Off On x: 0.0 y: 0.0 z: 1.0 r: 1.0 g: 1.0 b: 1.0
Reset to defaults
Apply

Figure 12: The lighting control dialog box.

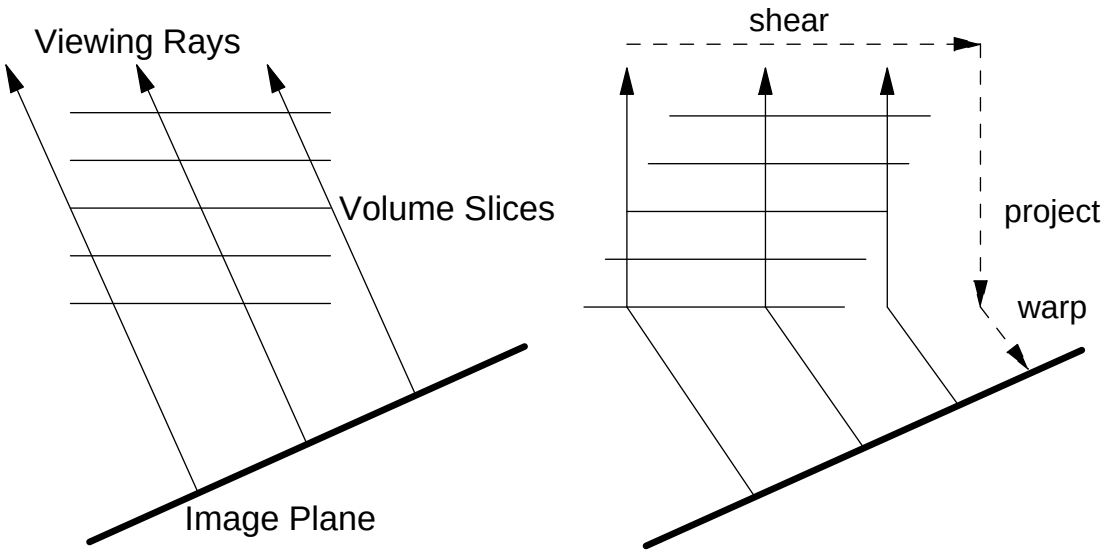

Figure 13: Volume rendering a 3D data block from an arbitrary viewing direction results in arbitrary angles of intersection between the viewing rays and the slices of the 3D data (a). Shear-warp volume rendering pre-calculates an overall shear of the volume that makes the viewing rays parallel to one axis of the sheared volume (b). The sheared data is then projected along the parallel rays into a buffer, which is finally warped onto the image plane to produce the rendered image. This results in a much more favourable alignment of the data in memory, vastly increasing rendering speed at the expense of a small reduction in the quality of the rendered image.

the direction of the light; values between -1.0 and 1.0 are valid. The final three components (r, g and b) specify the colour of the light as normalised RGB components; values between 0.0 and 1.0 are valid.

The default is to have a single white light shining down the z-axis; the "Reset to defaults" button resets all parameters to these default values. Changes to the lighting parameters are not implemented until the "Apply" button at the bottom of the dialog box is pressed, and even then have no effect unless the volume renderer is running. If the "Apply" button is pressed with invalid parameters in the dialog box, error messages will be printed to the text window of the top-level tinaTool and any lights with invalid parameters will be disabled. Finally, the lighting control dialog features an emergency mode to ensure that, even if all lights are disabled due to invalid parameters, the default lighting will be applied so that a rendering is still produced.

Note that, if coloured lights are specified with the renderer in greyscale mode, then the rendering will show the intensity of the lighting but not the colour. Switch to full colour mode to display the coloured lighting.

6.2 Limitations of the Volume Renderer

In order for the Manual Landmark tool to be usable, mouse interaction with the 3D image must be fast enough that the image can be manipulated in interactive time i.e. the renderer must be able to produce several frames per second. This places severe constraints on the amount of processing that can be performed during rendering. Traditionally, landmark identification software has used surface rendering algorithms in order to meet this requirement. Such algorithms rely on identifying a set of points on the bone surface within the volume, tesselating these points to produce a representation of the surface, and then only rendering this surface i.e. only rendering a small fraction of the data. However, this approach has a drawback in that it relies on a specific interpretation of the data i.e. some technique must be used to identify the surface prior to rendering. If, for example, a simple threshold is used, then this is equivalent to the assumption that all bone surfaces in the volume have a well-defined grey-level. At any point where this assumption fails (e.g. where the bone is thin and so obscured by partial voluming, is of low density, or is corrupted by image noise) errors will be generated in the rendering e.g. pseudofenestrations (holes in the rendered bone surface that do not exist in the original data).

The TINA Manual Landmark tool uses a volume rendering algorithm in order to avoid such problems. Volume rendering produces an image using the entire 3D volume: each pixel in the rendered image represents a vector through the 3D volume, and the renderer runs along these vectors, converting the intensity and intensity gradient of every voxel in the 3D volume into opacity and reflectivity, and combining these together to produce the intensity of the final pixel in the rendered image. The final, rendered image thus provides a more complete view of the data than a surface rendering, at the expense of significantly more processing.

In order to produce volume-rendered images in interactive time, the TINA Manual Landmark tool uses a fast volume rendering algorithm known as the shear-warp algorithm. The operation of this algorithm is illustrated in Fig. 13. Projecting along vectors from some arbitrary viewpoint would involve complex geometrical calculations in order to find which voxels were intersected by which vectors. Therefore, the shear warp algorithm calculates an overall shearing of the volume that makes all of the vectors parallel to the nearest major axis of the volume. The data is then projected into an intermediate image, which is parallel to one face of the volume, before being warped to produce the final rendered image. This results in a much more favourable alignment of the data in memory, vastly increasing rendering speed at the expense of a small reduction in image quality (introduced by the warping).

Several additional limitations have been implemented in order to maintain rendering speed. The renderer uses a parallel projection, rather than a perspective projection i.e. assumes that all rays from the viewpoint through the data are parallel, rather than diverging. The drawback of this approach is that it can introduce an optical illusion known as the Necker Cube [10] in which, particularly if the image has high transparency, it can be difficult to tell which side of the volume is closest to the viewer. Using renderer settings that minimise the translucency of bone surfaces will minimise this effect.

Furthermore, the lighting model used by the renderer does not take account of radiosity i.e. diffuse lighting reflected from multiple surfaces within the volume. Instead, it only takes account of direct rays of light from the light sources, reflected off of one surface to the viewer. This can cause a characteristic artefact in which rough bone surfaces appear to be speckled with small, dark patches. Correcting this with a more realistic lighting model would increase the processor time requirements to the point where the renderer could not run in interactive time. However, the effect can be minimised with careful choice of lighting parameters (e.g. using a single light shining onto the front of a rodent skull seems to give a good result).

Finally, the TINA Manual Landmark tool was designed primarily for use on very large micro-CT image volumes of rodent skulls (typically several GB in size). In order to make proper use of the accelerated 3D graphics hardware found in most PCs, the data to be rendered must be copied to the video card memory (in order to avoid continual paging of data from system memory to video card memory across the system bus). However, the majority of current video cards lack sufficient memory to hold an entire micro-CT rodent skull volume. Therefore, the volume renderer used in the Manual Landmark tool does not use graphics hardware; it runs on the CPU and from system memory.

7 The Landmark Points Dialog Box

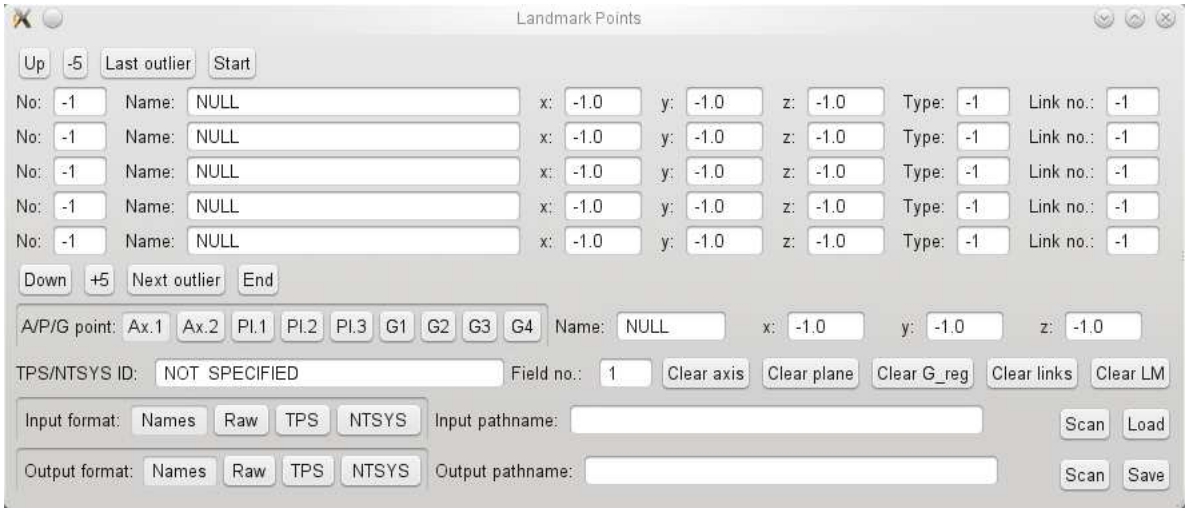
Landmark Points
Up
-5
Last outlier
Start
No: -1 Name: NULL x: -1.0 y: -1.0 z: -1.0 Type: -1 Link no.: -1
No: -1 Name: NULL x: -1.0 y: -1.0 z: -1.0 Type: -1 Link no.: -1
No: -1 Name: NULL x: -1.0 y: -1.0 z: -1.0 Type: -1 Link no.: -1
No: -1 Name: NULL x: -1.0 y: -1.0 z: -1.0 Type: -1 Link no.: -1
No: -1 Name: NULL x: -1.0 y: -1.0 z: -1.0 Type: -1 Link no.: -1
Down
+5
Next outlier
End
A/P/G point: Ax.1 Ax.2 Pl.1 Pl.2 Pl.3 G1 G2 G3 G4 Name: NULL x: -1.0 y: -1.0 z: -1.0
TPS/NTSYS ID: NOT SPECIFIED Field no.: 1 Clear axis Clear plane Clear G_reg Clear links Clear LM
Input format: Names Raw TPS NTSYS Input pathname: Scan Load
Output format: Names Raw TPS NTSYS Output pathname: Scan Save

Figure 14: The Landmark Points dialog box.

The TINA Manual Landmark tool was developed to allow the generation of lists of landmark point positions for use in Procrustes analysis. In this type of analysis, it is essential that each list of landmark points contains the same set of physiological locations in the same order. Therefore, rather than allowing the free-form generation of lists of points, the Manual Landmark tool requires the user to prepare a text file containing a list of point names and point numbers, referred to as a "names file", which must be loaded into the tool before point locations can be specified. An example is shown below. The file can be prepared in any text editor, and on each line should contain a point number at the start of the line, followed by whitespace (one or more spaces or tab characters), then the text description of the point.

In many geometric morphometric studies, some or all of the landmarks will occur in pairs arranged symmetrically on either side of the plane of bilateral symmetry. The "Linklines" option in the 3D TV (see Section 5.2.2) displays lines linking these pairs of points. This provides a quick, visual check at the end of the manual landmarking process; the lines should approximately be parallel and, if this is not the case, it may indicate that the landmarks have not been identified in the correct order e.g. a pair of points has been transposed. In order to use this option, a file listing the links must be prepared. Each line of the file contains two landmark point numbers, corresponding to the landmark point numbers in the names file, separated by whitespace (spaces or tab characters). The file must have the same filename as the landmark names file and must have the extension .llf (e.g. if the names file is called points.txt, then the linking lines file must be called points.llf). It must also be placed in the same directory. The linking lines file will then be automatically loaded or saved whenever the names file is loaded or saved. Lines will then be drawn in the 3D TV between the specified points whenever the "Linklines" option is selected in the "3D Tv draw" check list and coordinates have been specified for both points. An example linklines file is shown below.

An example landmark point names file for mouse skulls:

1 Most anterior point of the anterior palatine foramen, left side.
2 Most anterior point of the anterior palatine foramen, right side.
3 Intersection of frontal process of maxilla with frontal and lacrimal bones, left side.
4 Intersection of frontal process of maxilla with frontal and lacrimal bones, right side.
5 Frontal-squasmosal intersection at temporal crest, left side.
6 Frontal-squasmosal intersection at temporal crest, right side.
7 Center of alveolar ridge over maxillary incisor, left side.
8 Center of alveolar ridge over maxillary incisor, right side.
9 Lateral intersection of maxilla and palatine bone posterior to the third molar, left side.
10 Lateral intersection of maxilla and palatine bone posterior to the third molar, right side.
11 Anterior notch on frontal process lateral to infraorbital fissure, left side.
12 Anterior notch on frontal process lateral to infraorbital fissure, right side.
13 Most posterior point of the anterior palatine foramen, left side.
14 Most posterior point of the anterior palatine foramen, right side.
15 Anterior-most point at intersection of premaxillae and nasal bones, left side.
16 Anterior-most point at intersection of premaxillae and nasal bones, right side.
17 Most infero lateral point on premaxilla-maxilla suture, left side.
18 Most infero lateral point on premaxilla-maxilla suture, right side.
19 Intersection of parietal, temporal and interparietal bones,left side.
20 Intersection of parietal, temporal and interparietal bones,right side.
21 Most inferior aspect of posterior tip of medial pterygoid process,left side.
22 Most inferior aspect of posterior tip of medial pterygoid process,right side.
23 Joining of squasmosal body to zygomatic process of squasmosal, left side.
24 Joining of squasmosal body to zygomatic process of squasmosal, right side.
25 Intersection of zygomatic process of maxilla with zygoma (jugal), superior surface, left side.
26 Intersection of zygomatic process of maxilla with zygoma (jugal), superior surface, right side.
27 Intersection of zygoma (jugal) with zygomatic process of temporal, superior aspect, left side.
28 Intersection of zygoma (jugal) with zygomatic process of temporal, superior aspect, right side.
29 Basion.
30 Bregma: intersection of frontal bones and parietal bones at midline.
31 Intersection of parietal bones with anterior aspect of interparietal bone at midline.
32 Intersection of interparietal bones with squamous portion of occipital bone at midline.
33 Nasion: Intersection of nasal bones, coudal point.
34 Nasale: Intersection of nasal bones, rostral point.
35 Opisthion.
36 Intradentale superior.
37 Most posterior projection of the posterior nasal spine.
38 Most superior point on the external auditory meatus, right side.
39 Anterior edge of the alveolar process where right first molar hits alveolus, right side.
40 Most anterior-medial point on the right carotid canal, right side.
41 Most superior point on the external auditory meatus, left side.
42 Anterior edge of the alveolar process where right first molar hits the alveolus, left side.
43 Most anterior-medial point on the left carotid canal, left side.
44 Anterior nasal spine.
45 Intersection of the right occipital condyle and the foramen magnum, taken at the lateral most curvature, right side.
46 Intersection of the left occipital condyle and the foramen magnum, taken at the lateral most curvature, left side.

An example landmark point links file for mouse skulls (corresponding to the names file above):

```
1       2
3       4
5       6
7       8
9      10
11     12
13     14
15     16
17     18
19     20
21     22
23     24
25     26
27     28
38     41
39     42
40     43
45     46
```

Names files and other files containing landmark points are loaded and saved using the Landmark Points dialog box, which also displays the currently loaded list of landmark points and provides some basic tools to manipulate it. This dialog box is started by clicking on the "Landmark List" button at the top of the Manual Landmark tool, and is shown in Fig. 14.

In order to load a landmark names file:

- Prepare the names file in advance, using any text editor. Optionally, prepare the corresponding linklines file.
- In the Landmark Points Dialog box, enter the absolute or relative pathname of the landmark names file into the "Input Pathname" field, either directly or by using the file-browser started by clicking on the "Scan" button.
- Select the format of the file (in this case "Names") from the "Input format" choice list.
- Click the "Load" button in the Landmark Points dialog box.

Once a names file has been loaded, the names and numbers of the landmarks can be viewed in the Landmark Points dialog box. The top half of the dialog box shows five landmarks at a time, displaying their number, name, any stored x, y and z coordinates, a "type" field (see below), and a link no. field, which specifies which (if any) other point the point is linked to when the "Linklines" option is used in the "3D Tv draw" checklist. The user is free to edit any of these fields manually; however, the type code has specific meaning that affects the operation of the software, particularly the Automatic Landmark Point Placement tool (see Section 7.1). For example, the Manual Landmark tool assigns a type of -1 to any point that has not yet been identified, and a type of 0 to any point that has been manually identified. Any point with a type of -1 will not be displayed in the 3D TV. Therefore, the point type code should not, in general, be altered by the user.

After loading a landmark names file, one landmark in the list is always considered to be the current landmark. During manual landmarking, whenever a point is identified the coordinates of the 3D cursor will be transferred to the current landmark. The current landmark is always the one displayed in the middle of the five rows of landmark information in the Landmark Points dialog box. The user can specify the current landmark (so that points can be marked up in any order) by scrolling up and down through the landmark list using the "Up" and "Down" buttons in the Landmark Points dialog box. These buttons are also assigned keyboard shortcuts in all TVs of the Manual Landmark tool (see Sections 5.2.1, 5.2.2 and 9). In addition, the "-5" and "+5" buttons move up and down by five points in the landmark list, and the "Start" and "End" buttons move to the first and last landmarks respectively. Finally, the "Last outlier" and "Next outlier" buttons move to the previous and subsequent points with type codes greater than 2 (if any) i.e. points that have been identified as potential errors in the automatic landmarking process (see Section 13.7). These buttons allow the user to scroll through the landmark list quickly when correcting any errors from the automatic landmark point placement software (see Section 13.8).

In order to avoid having to keep the Landmark Points dialog box open during manual landmark identification, a

subset of the dialog is also available in the Manual Landmark tool, which displays the number, name and type code of the current landmark and has copies of the "Up", "Down", "-5", "+5", "Start", "End", "Last outlier" and "Next outlier" buttons. However, the current landmark is always kept consistent between these two windows, so when scrolling through the landmark list using the buttons in the Manual Landmark tool, the user will also see the list scroll in the Landmark Points dialog box.

The Landmark Points dialog box can also load and save a variety of other file formats, as specified in the "Input format" and "Output format" choice lists. "Raw" is a TINA-specific, ASCII text format that contains all of the information from the landmark list. "TPS" and "NTSYS" are standard formats for morphological data. All of these formats are ASCII text files, and so can be viewed or edited in any text editor.

In order to save a landmark list:

- In the Landmark Points dialog box, enter the absolute or relative pathname of the landmark file into the "Output pathname" field, either directly or by using the file-browser started by clicking on the "Scan" button.
- Select the format of the file from the "Output format" choice list.
  - Raw files can only hold one list of landmark points. If a Raw file is overwritten, then the previous version of the file will be replaced
  - TPS files can hold multiple lists of landmark points. If landmark points are saved to a TPS file that already exists, then the new set of landmark points will be appended to the existing file. The TPS format assumes that each set of landmarks in the file has a unique identifier: the contents of the "TPS/NTSYS ID" field will be used as the identifier. The software does not check whether this is unique in the file i.e. it is up to the user to do this. However, the name of the directory containing the image files is copied to the "TPS/NTSYS ID" field whenever a sequence is loaded and so, if the image volumes for individual specimens are stored in individual directories with meaningful names (reflecting the name or code for the specimen) then this process will be handled automatically.
  - NTSYS files can hold multiple lists of landmark points. If landmark points are saved to a NTSYS file that already exists, then the new set of landmark points will be appended to the existing file. The contents of the "TPS/NTSYS ID" field will be saved in the file as a comment. The landmark point names will be saved to file as a list immediately preceeding the data matrix; however, NTSYS only supports labels of up to 16 characters, and so the names will be abbreviated if necessary. Users should be aware of this when preparing lists of landmark point names if they intend to use the NTSYS file format. The landmark point coordinates will be saved to the data matrix of the NTSYS file as columns (i.e. each column of the matrix represents a point).
- Click the "Save" button in the Landmark Points dialog box.

Landmark point information can also be loaded from any of these file formats. The "Field no." field allows the user to specify which record to load from a TPS or NTSYS file that contains more than one set of landmark points. Since the software can only maintain one list of landmark points at a time, a conflict may arise in situations where the user attempts to load a file when some landmark information has already been loaded. In such situations, the lengths of the two lists will be compared. If they are the same, then the software will assume that the points are in the same order and will load all of the information contained in the new file into the relevant fields of the current landmark list. If the lengths are different, then the software will either refuse to load the new file, or will discard the current landmark list and replace it with the information from the file, always choosing the option that discards the least data i.e. if the user attempts to load a names file after a TPS or Raw file has been loaded, the software will refuse to load the file; if the user attempts to load a TPS or Raw file after loading a names file, then the names will be replaced with those from the file. However, the Landmark Points dialog box provides a "Clear Landmarks" button that allows the current landmark list to be cleared, and also provides "Clear Axis", "Clear Plane", "Clear links" and "Clear G-reg" buttons to allow the user to discard any currently stored axis point, plane points (see Section 8.1), linklines information, or G-reg points (see Section 13.2.1).

In order to load a landmark list:

- In the Landmark Points Dialog box, enter the absolute or relative pathname of the landmark file into the "Input pathname" field, either directly or by using the file-browser started by clicking on the "Scan" button.
- Select the format of the file from the "Input format" choice list.
- If loading an NTSYS or TPS file containing multiple records, enter the ordinal position of the record in the file into the "Field no.:" field.
- Click the "Load" button in the Landmark Points dialog box.
  - If a record from a TPS file is loaded, the ID of the record will be copied to the "TPS/NTSYS ID" field.
  - If a record from an NTSYS file is loaded, any comments in that record will be displayed in the text window at the bottom of the top-level tinaTool window.

Note that Raw, NTSYS and TPS files do not contain the information from the line linking lines file. However, the user can load this information by loading a names file (with a corresponding linking lines file) after loading a Raw, NTSYS or TPS file, as long as all files contain corresponding lists of landmark points. Note also that the Raw format contains more information that NTSYS or TPS. In particular, it records the landmark type code (see Sec. 7.1), which is used by the Automatic Landmark Point Placement tool (see Sec. 12) to record information about errors in the automatic point locations. Therefore, Raw is the preferred format for saving results from automatic landmark point location if the user intends to check and correct outliers at a later time.

7.1 Landmark Type Codes

Each landmark, axis, plane or global registration (G-reg) point stored in the software is assigned a type code, which can be seen in the Landmark Points dialog box. These are used by the software to represent information about the point and should not, in general, be changed by the user.

Point type codes:

-1	NULL point (not yet allocated).
0	Point loaded from TPS/NTSYS/names file.
1	Point with manually specified coordinates.
2	Point with automatically identified coordinates.
3-9	Automatically identified point flagged as an outlier.
101	First axis point.
102	Second axis point.
103	First plane point.
104	Second plane point.
105	Third plane point.
106	First global registration point.
107	Second global registration point.
108	Third global registration point.
109	Fourth global registration point.

8 Manual Landmark Identification

The aim of the Manual Landmark tool is to allow the user to manually specify morphological landmarks in any 3D medical image volume that has been loaded into TINA via the Sequence tool. In order to do this, the user must manipulate the position of the 3D cursor until it corresponds to the desired point, and then store the coordinates of the 3D cursor. When data is first loaded, the 3D cursor if initialised at the centre of the volume. Its position is displayed graphically in all TVs as the red cross-hair. Its coordinates are also displayed numerically in the X, Y and Z fields of the Manual Landmark tool.

The position of the 3D cursor can be manipulated in a number of ways:

- Numerical values can be entered directly into the X, Y and Z fields of the Manual Landmark tool.
- The arrow keys after each of the X, Y and Z fields in the Manual Landmark tool allow the user to increment or decrement each coordinate of the 3D cursor: the size of the increment can be chosen via the "Resolution" choice list in the Manual Landmark tool.
- When the 2D TV mouse interaction is in "Pick" mode, the user can move the large red cross-hair representing the 3D cursor by left-clicking on a position in any 2D TV, or by left-clicking and dragging. Alternatively, the keyboard shortcuts can be used to shift the position of the 3D cursor (see Section 5.2.1).
- When the 3D TV mouse interaction is in "Pick" mode, the user can left click on any position in the 3D TV: the 3D cursor will then be moved to the upper-most bone surface under the mouse cursor. The bone surface is detected using a simple intensity threshold, specified in the "Bone Threshold" field of the Manual Landmark tool (see Section 5.2.2).

Once the 3D cursor is in the desired position, its coordinates can be stored by clicking the "Mark Point" button in the Manual Landmark tool. The coordinates will be transferred to the point specified in the "Markup" choice list. If "Curr. LM" is selected, then the coordinates will be transferred to the current landmark i.e. the landmark specified in the "No." and "Name" fields in the Manual Landmark tool, and also displayed in the middle of the five rows of landmark information in the Landmark Points dialog box. A list of landmark names must already have been loaded in order to do this: otherwise, the software will issue a warning that no landmark list has been loaded. After transferring the coordinates, the software will also increment the current landmark (i.e. move to the next one in the landmark list), so that the user does not need to scroll through the landmark list during landmarking. The other choices in the "Markup" list allow the user to specify axis, plane and G-reg points (see Sections 8.1 and 13.2.1).

The "Jump lock" switch and the "Jump stored" button allow the user to re-use coordinates of landmarks that have previously been identified. If the user scrolls through the landmark list (using e.g. the "Up" and "Down" buttons in the Manual Landmark tool or in the Landmark Points dialog box) to change the current landmark to a landmark for which coordinates have already been stored, and then clicks the "Jump stored" button, the 3D cursor will be moved to the stored coordinates for the current landmark. If the "Jump lock" switch is set to "On", then this becomes the default behaviour when scrolling through the landmark list: the 3D cursor will automatically be moved to the stored coordinates of the current landmark whenever such stored coordinates exist.

The "Snap" button provides a basic snap-to-grid style functionality. If the 3D cursor is not located on a bone surface, as specified by the "Bone threshold" field of the Manual Landmark tool, then clicking on this button will move it to the nearest bone surface.

8.1 Axis and Plane Points

In addition to landmark points, the TINA Manual Landmark tool supports three plane points and two axis points. The plane points can be used, for example, to specify the plane of bilateral symmetry. The axis points are used to specify an arbitrary axis within the 3D image. Both sets of points are used to implement various utility functions that make the manual landmarking process easier. Axis and plane points can be specified regardless of whether a landmark list is loaded or not.

In order to mark up an axis or plane point:

- Move the 3D cursor to the desired point in the 3D image;
- Select the desired point in the "Markup" list of the Manual Landmark tool (e.g. "Pl.1" is the first plane point);
- Press the "Mark Point" button in the Manual Landmark tool (or use the keyboard shortcut assigned to this function in any of the Manual Landmark tool TVs; see Sections 5.2.1 and 5.2.2): this will save the current coordinates of the 3D cursor to the point specified in the "Markup" list of the Manual Landmark tool.

In order to re-use an existing landmark point as an axis or plane point:

- Scroll through the list of landmark points, either in the Landmark Points dialog box or in the Manual Landmark tool.
- Press the "Jump to" button in the Manual Landmark tool: this will move the 3D cursor to the coordinates of the current landmark point, as specified in the Landmark Points dialog box and in the Manual Landmark tool.
- Select the desired point in the "Markup" list of the Manual Landmark tool (e.g. "Pl.1" is the first plane point).
- Press the "Mark Point" button in the Manual Landmark tool: this will save the current coordinates of the 3D cursor to the point specified in the "Markup" list of the Manual Landmark tool

In order to display axis or plane points within the 3D rendering, select the desired options in the "3D Tv Draw:" check list in the Manual Landmark tool. Plane points are displayed in cyan, and axis points in yellow. As with landmark points and the 3D cursor, the points are displayed as 3D cross-hairs by default, but can be displayed as spheres if the "Ball" option is selected in the "3D Tv Draw:" check list. Any parts of a cross-hair or sphere that lie behind a bone surface (as specified by the "Bone threshold" field in the Manual Landmark tool) from the current viewpoint are displayed in a lower intensity, whilst points that are not behind a bone surface are displayed in a higher intensity; this provides 3D cues to the user, making it easier to see how each point interacts with the bone surface. However, this functionality is disabled during rotation of the 3D rendered image in order to prevent slowdown; during rotation, all points are displayed in a medium intensity. Additionally, selecting the "Poly" option in the "3D Tv Draw:" check list will display a yellow line representing the axis, if both axis points have been defined and the "Axis" option is on, and a cyan grid representing the plane, if all three plane points have been defined and the "Plane" option is on.

The stored coordinates for axis and plane points can be viewed in the Landmark Points dialog box. The "A/P/G point:" choice list selects which point will be viewed, and the "Name:", "x:", "y:" and "z:" in the same row of the tool display the name and coordinates of the selected point (if stored). Any current axis and plane points can be deleted using the "Clear Axis" and "Clear Plane" buttons in the Landmark Points dialog box.

Plane points are used by the "Reflect cursor" button in the Manual Landmark tool. Once three plane points have been specified, pressing this button will project the current 3D cursor position through the plane, to an equidistant point on the other side, and reset the 3D cursor to that point. This can be useful when marking up symmetrical structures.

Axis points are used to limit the rotation of the volume rendered image in the 3D TV. Once both axis points have been specified, press the "Lock Rotation" button in the Manual Landmark tool. The volume rendered image will then only rotate around an axis parallel to that defined by the two axis points and passing through the centre of the volume. This functionality can be useful in viewing the extremal extents of some structures.

The viewpoint of the 3D TV can also be aligned using the "Rotate to Axis" and "Rotate to Plane" buttons in the Manual Landmarking tool. Once three plane points have been selected, clicking the "Rotate to Plane" button will rotate the volume such that the viewing direction after rotation is the normal to the plane, using the smallest rotation angle possible. Clicking the button a second time will flip the volume over to give the view from the other side of the plane. Once all axis and plane points have been defined, clicking the "Rotate to Axis" button will rotate the volume such that the projection of the axis onto the plane is aligned with the left-right axis of the 3D TV, so that using both buttons in combination will completely align the volume. Again, the rotation applied is through the smallest angle possible, and clicking the button a second time will flip the volume over to give the view from the opposite side.

9 Keyboard Controls

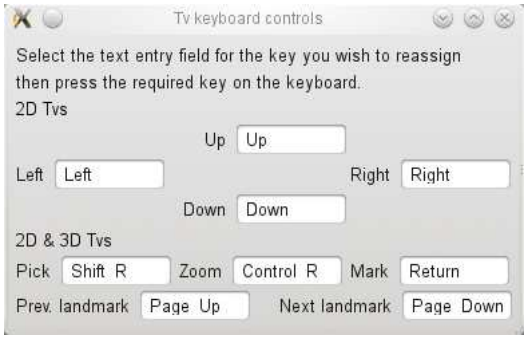

Figure 15: The TV Keyboard Controls dialog box.

As described in Sections 5.2.1 and 5.2.2, a number of keyboard shortcuts are available in the 2D and 3D TVs, enabling commonly used functions to be applied without moving the mouse cursor away from the TV tools. These can be reassigned to any key (with the exception of a few keys reserved by the window manager, such as the print screen key) using the TV Keyboard Controls dialog box. This dialog box is started by clicking the "Keyboard Controls" button at the top of the Manual Landmark tool, and is shown in Fig. 15. In order to reassign a key, left-click with the mouse inside the field displaying the keyboard control you want to reassign, and then press the key you want to specify. The name of the key will appear in the entry field. The same set of keys are used across all TVs.

10 The View Tool

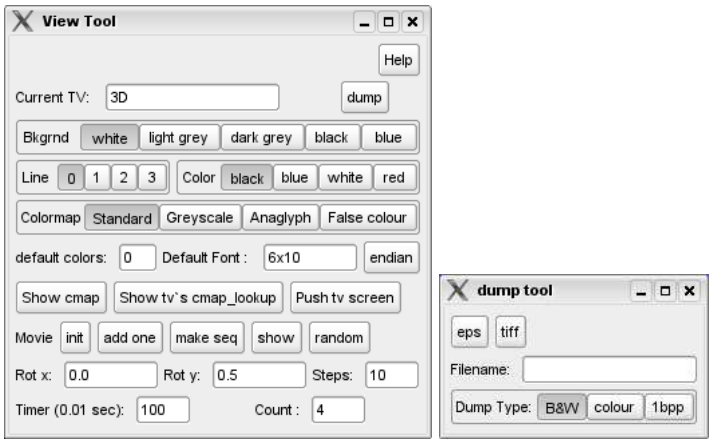

Figure 16: The View tool and Dump tool.

The View and Dump tools provide additional interaction with graphics displayed in TV tools. See the TINA User's Guide [19] for full instructions on how to use these tools. Two functions of these tools are of particular interest when using the TINA Manual Landmark tool: the creation of movies from 3D images and the saving of images to files.

The View tool can be used to prepare a movie from any TV displaying a 3D image (e.g. the 3D TV of the Manual Landmark tool), in which the image rotates in 3D around its central point. Each frame of the movie is rendered in turn and stored as a 2D image within the TV tool; during playback, these 2D images are displayed in the TV as a slide show. Therefore, no rendering is performed during playback, allowing the movie to play in a much shorter time than it took to render. This can significantly increase the viewer's perception of 3D, and so can be useful in seeing exactly where landmark points are located in the 3D image. The user can specify the angles of rotation, the number of frames in the movie, the time delay between each frame during playback, and the number of times the movie loops during playback. Note that the rotation angles are specified in radians, around the current x-axis (i.e. left-right axis) and y-axis (i.e. up-down axis) of the TV i.e. they do not refer to the coordinate system used in the Manual Landmark tool. Once a movie has been created, it is stored in the TV tool until the user deletes it.

To create a movie in a TV tool displaying a 3D image:

- Select the TV for which you wish to produce a movie by clicking on its entry in the TV list of the tool it is associated with (e.g. click on "3D" in the TV list of the Manual Landmark tool).
- Start the View tool by clicking the "View" button in the tinaTool window. The View tool will start, and will display the name of the currently selected TV in the "Current Tv" field.
- Select the parameters of the movie:
  - Rot x: the rotation to apply about the x-axis of the image (in radians).
  - Rot y: the rotation to apply about the y-axis of the image (in radians).
  - Steps: the number of frames in the movie.
  - Timer (0.01 sec): the time delay between each frame during playback.
  - Count: the number of times to loop through the movie.
- Click "Make seq" to prepare the movie and store it to the TV.
- Click "Show" to play the movie in the TV.
- Clicking "Init" will delete any previously stored movie, allowing a new one to be prepared.

The Dump tool, which is a sub-tool of the View tool, can save the image displayed in any TINA TV tool as a TIFF or Encapsulated Postscript file; TIFF files can be used to transfer images to Windows machines e.g. for incorporation into MS Word documents, whilst EPS files can be used to incorporate images into documents prepared under Linux (e.g. using Latex).

To save the image displayed in a TV tool to a file:

- Select the TV for which you wish to produce a movie by clicking on its entry in the TV list of the tool it is associated with (e.g. click on "3D" in the TV list of the Manual Landmark tool).
- Start the View tool by clicking the "View" button in the tinaTool window. The View tool will start, and will display the name of the currently selected TV in the "Current Tv" field.
- Click the "Dump" button in the View tool. This will start the Dump tool. Enter a relative or absolute pathname (relative pathnames are relative to the directory containing the tinaTool executable) into the "Filename" field, and select the colour depth using the "Dump Type" field: "B&W" will give a greyscale image, "Colour" will give a full-colour image, and "1bpp" will give a black and white image (which can be useful for saving graphs displayed in the Imcalc Graph TV).
- Click on either the "TIFF" or "EPS" button to save the image from the currently selected TV to file in TIFF or Encapsulated Postscript format.

11 Using the Manual Landmark Tool

A complete set of steps required to specify a set of landmarks is as follows:

- Prepare a landmark names file, specifying the points that you intend to identify (see Section 7). Optionally, prepare the linklines file.
- Start TINA; start the Manual Landmark tool.
- Start the Sequence tool, assign a TV tool to the sequence tool TV (optional; see Section 3). and load a 3D medical image volume (see Section 4).
- Start the volume renderer, using the "3D Tv" switch in the Manual Landmark tool.
- Start four new TV tools, and install them on the four TVs of the Manual Landmark tool.
- Start the VR Control dialog box and the Lighting Control dialog box, and manipulate the volume renderer options to produce a good volume rendering (see Sections 6 and 6.1 for suggested settings).
- Use the techniques described in Section 5.2.2 to specify a suitable bone threshold.
- Load the landmark names file using the Landmark Points dialog box (see Section 7).
- The current landmark will now be the first one in the list: manipulate the rotation of the image in the 3D TV (with the 3D TV mouse interaction in "Zoom" mode) until you can see the corresponding point on the bone surface.
- Switch the 3D TV mouse interaction to "Pick" mode, and then left-click on the position of the landmark point in the 3D TV: the 3D cursor will be moved to that point.
- Check the position of the 3D cursor in the 2D TVs: if necessary, refine it using the arrow keys in the Manual Landmark tool, or by using mouse or keyboard interaction with the images in the 2D TVs (see Section 5.2.1).
- When you are happy with the position of the 3D cursor, click the "Mark point" button in the Manual Landmark tool (or use the keyboard shortcut in the 2D or 3D TVs: see Sections 5.2.1 and 5.2.2). The coordinates of the 3D cursor will then be transferred to the current landmark, and the current landmark will then be incremented (i.e. the next landmark in the list will become the current landmark).
- Proceed through the landmark list, marking up each point in turn until the coordinates of all points have been specified.
- Save the landmark list to file using the Landmark Points dialog box (see Section 7).
- At this point the user can load another 3D image volume and repeat the process. Note that the previous landmark list will still be in memory, and so can be used as starting points (if the data sets are similar enough) or can be deleted using the "Clear Landmarks" button in the Landmark Points dialog box, allowing the names list to be re-loaded.

Part 2: The Automatic Landmark Point Placement Tool

P.A. Bromiley
Imaging Science and Biomedical Engineering Division,
Medical School, University of Manchester,
Stopford Building, Oxford Road,
Manchester, M13 9PT.

12 Introduction

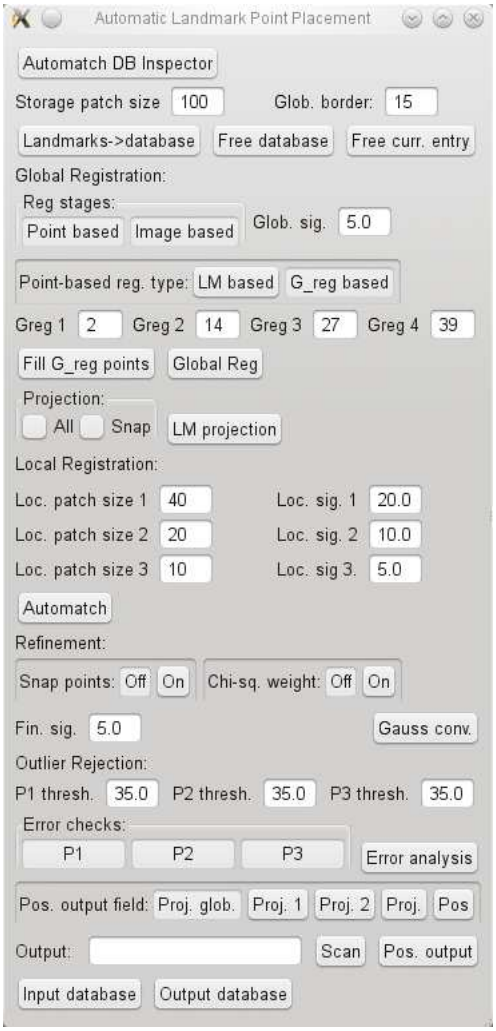

Figure 17: The Automatic Landmark Point Placement tool.

In order to accelerate the process of landmark identification, the TINA Geometric Morphometrics toolkit is capable of automatically estimating landmark positions, based on a database of previous examples of manually identified landmarks. This functionality is available via the Automatic Landmark Point Placement tool, which can be launched by clicking on the "Automatch" button in the Manual Landmark tool. The Automatic Landmark Point Placement tool is shown in Fig. 17.

The automatic landmark identification operates by storing information from manually identified sets of landmarks, including both the coordinates of the landmark points and image data from around the points, in a database (the "automatch database"). These image patches are then aligned with the data set currently loaded into the Sequence tool, using a coarse-to-fine registration process, producing multiple estimates of the position of each landmark (one for each entry in the database). The registration is based on a maximum likelihood estimation of a nine-parameter affine transformation model i.e. a 3D translation, rotation and scaling. Rather than using 3D blocks of data around each landmark point, the software extracts 2D images on the three orthogonal planes intersecting at the position of the landmark point and aligned with the major axes of the image volume; this provides enough data to estimate the parameters accurately and reduces the processor time required compared to the use of 3D data blocks. Smoothing is applied to the patches in order to reduce the effects of noise, and spatial derivatives are taken in the x and y directions of each patch, producing a total of six image patches for each landmark point. The likelihood function itself is the sum-squared difference between the data in these patches and the data in the image volume currently loaded into the Sequence tool i.e.

$$lnL = \sum_{i=0}^{n} \frac{\gamma I_i - T(J)_i}{\sigma_e} \tag{1}$$

where L is the likelihood, I represents smoothed image derivatives of the patches stored in the database, J represents smoothed image derivatives of patches from the volume currently loaded into the sequence tool, T represents the nine-parameter affine transformation model being optimised in order to align the two sets of patches, n represents iteration over corresponding voxels in the image patches, γ is a relative scaling between the I and J patches, and σ_e is the standard deviation of the noise on the numerator of the likelihood function. The noise is calculated by measuring the noise on the original I and J image volumes using the standard TINA noise measurement function, which is based on measuring the width of zero-crossings in vertical and horizontal gradient histograms [13], and then applying error propagation to account for the effects of smoothing, differentiation, and multiplication by γ. The use of image derivatives, rather than intensities, makes the software robust to the presence of any overall offset between the intensities in the two image volumes. The γ factor provides further robustness against any linear scaling between the intensities in the two volumes. Since γ must be taken into account in the calculation of σ_e, and the noise estimate cannot be allowed to vary during maximum likelihood optimisation, γ is estimated once, at the end of the first, global registration stage (see below) and then fixed during local registration. The optimisation itself is performed using the simplex algorithm [14]. Accurate knowledge of n and σ_e also allows the calculation of a goodness-of-fit score from the final value of the likelihood function, as the likelihood can be translated directly into a χ^2 using

$$-2lnL = \chi^2 \tag{2}$$

(the χ^2 per degree of freedom can then be obtained by dividing by the number of independent terms in the likelihood sum minus the number of parameters optimised).

Since the automatic point location process is based on affine transformations i.e. assumes that an affine transformation of the image volume currently loaded into the Sequence tool provides a good model of the image volumes stored in the database, it is important that the current and stored volumes have similar shapes e.g. for identification of morphometric landmarks on the bone surface of a given sample, testing on micro-CT images of rodent skulls has indicated that the software is reliable if the database contains examples from the same genus as the current volume. The list of landmark points must also be identical i.e. the must contain the same number of points, marked up in the same locations and in the same order. The software will not allow the addition of new entries to the database if the length of the list of landmark points differs from previous entries already stored in the database, but it is up to the user to ensure that the lists of landmarks contain the same biological points marked in the same order.

Since the aim of the software is to facilitate the placement of landmark points for experiments in geometric morphometrics, there will naturally be some difference in shape between the data sets being analysed. Therefore, the software uses a coarse-to-fine registration strategy, in which all stages are based on the likelihood function given above, in order to progressively align the two sets of image patches. The first stage is a global registration, in which the patches span the entire volume and intersect at its centre. The shape difference between the structures in the images in the current volume and the database mean that the global patches from the former are a relatively poor model of the latter, and so the result is an approximate alignment. The registration then proceeds to align patches around individual landmark points. Three stages of this local registration occur with progressively smaller patches each time. This is based on the concept that, as the patch size is decreased, the effect of shape differences between the database and current volumes is reduced i.e. the latter becomes a progressively better model of the former. The global and initial local stages of registration therefore search the image volume to find approximate positions, which are refined in the later stages of local registration to improve accuracy.

13 Use of the Software

Automatic landmark identification involves the following stages:

- Database construction;
- Initial, point-based global alignment;
- Image-based refinement of global alignment;
- Coarse-to-fine local alignment;
- Final refinement of landmark estimates;
- Outlier detection.

These are described in detail in the following sections.

13.1 Database construction

Automatic landmark identification is based on a database of examples of manually identified landmarks, which must be constructed by the user. The database should contain typically around eight to twelve sets of manually identified landmarks. In order to construct the database, load an image volume and manually identify the list of landmarks (or load them from file if available). The user must also consider the initial, point-based stage of registration (see Section 13.2.1), and so may wish to mark the four G-reg points. Once the coordinates of all landmarks have been manually identified, press the "Landmarks->database" button in the Automatch Dialog box. This will extract all of the information required for subsequent operation of the automatic landmark software and enter it into the database.

All image-based registration stages allow the application of smoothing during the preparation of the image patches, in order to reduce image noise and thus provide a smoother cost function for the optimisation, reducing the probability that the optimisation will get trapped in a local minimum. In order to allow accurate calculation of the smoothed images, the software supports the use of a border around the patches i.e. it stores and smooths a patch of the requested patch size plus the border size, and then ignores the data in the border region during calculation of the likelihood. The border size should be set to at least three times the size of the smoothing kernel in order to ensure that no part of the smoothing kernel falls outside the patch when calculating any smoothed image intensity that will be used in the calculation of the likelihood (since the smoothing function truncates the smoothing kernel at the 3σ points). Borders/patch sizes are controlled by two parameters: "Storage patch size" and "Glob. border". The first controls the size of the images patches around each landmark point that are stored in the database. The patches must be large enough to include the largest patch size used in the local registration plus a border of three times the size of the smoothing kernel. For example, if the coarsest stage of local registration uses a patch size of 40 and a smoothing kernel size of 20, then the stored patch size must be at least $40 + 3 \times 20 = 100$. The software will abort the local registration if the stored patch size is too small. Similarly, "Glob. border" controls the size of the border around the full-sized patches through the centre of the volume that are used in the image-based stage of global registration. This should be at least three times the size of the smoothing kernel that will be used in the global registration. The software will abort the global registration if the border is too small. The user must take this into consideration if they wish to experiment with different settings of the smoothing kernel and local registration patch size parameters. The patch, border and smoothing kernel sizes are all measured in the coordinate system of the current volume i.e. taking into account any down-sampling applied during data loading. The patch size parameters refer to the distance from the centre of the patch to the edge i.e. if a patch size parameter is set to 10, the patch used measures 20×20 voxels in the down-sampled coordinate system.

The database can be stored to disk by entering the desired path and filename (pathnames can be either relative or absolute) into the "Output" field of the Automatch Dialog box, and then clicking on the "Output AM Database" button. Note that database files are in a binary format and cannot be manually edited. In order to load a previously prepared database, enter the path and filename into the "Output" field and click the "Input AM Database" button; any database currently loaded into the software will be replaced. The database currently loaded into the software can be freed by clicking on the "Free Database" button. The contents of the database can be examined using the Automatch Database Inspector dialog box (see Section 17). Additionally, the current entry in the automatch database (i.e. the one displayed in the Automatch Database Inspector dialog box) can be deleted by clicking the "Free curr. entry" button in the Automatic Landmark Point Placement tool; this allows correction of errors during database construction.

13.2 Global Registration

The global registration consists of two stages; an initial, landmark point based stage and a final, image-based stage. These are described in detail in the following sections. The "Reg stages" check list determines which stages will be run; in general, both "Point based" and "Image based" should be selected (however, see Section 15 for an alternative way to operate the software). Global registration is performed by clicking on the "Global Reg" button; the software will then proceed to optimise transformation models that bring the image volume currently loaded into the sequence tool into alignment with the entries in the automatch database.

13.2.1 Initial Point-based Global Registration

Image-based registration algorithms cannot reliably handle gross misalignments (for example, if the current volume is rotated by approximately 180 degrees compared to those stored in the database). Therefore an initial, approximate stage of alignment is required prior to image-based registration. In the current version of the software (see Section 18), this requires the user to mark four landmark points in each volume; the first stage of the global registration is based on aligning these points by minimising the sum-squared distances between corresponding points using the simplex optimisation algorithm. These points must span all three dimensions of the space and be well separated in order to provide an accurate alignment i.e. they must not be co-planar. The points can be landmark points i.e. points in the current landmark list; simply mark up any four or more of these and the software will calculate the initial global registration by aligning them to the corresponding points in the database. Alternatively, the user may wish to use points not included in the landmark point list for the global alignment. The four G-reg points provide this functionality; they can be manually identified in the same way as any other landmark point, using the "G1", "G2", "G3" and "G4" selections in the markup choice list. The "3D Tv draw" check list in the Manual Landmark tool provides the option to display the positions of the G-reg points in the 3D TV, where they are displayed as orange spheres or cross-hairs. The stored coordinates for G-reg points can be viewed in the Landmark Points dialog box. The "A/P/G point:" choice list selects which point will be viewed, and the "Name:", "x:", "y:" and "z:" fields in the same row of the tool display the name and coordinates of the selected point (if stored).

In order to identify a G-reg point:

- Move the 3D cursor to the desired point in the 3D image.
- Select the desired point in the "Markup" list of the Manual Landmark tool (e.g. "G1" is the first G-reg point).
- Press the "Mark Point" button in the Manual Landmark tool: this will save the current coordinates of the 3D cursor to the point specified in the "Markup" list of the Manual Landmark tool.

The choice of whether to use the landmark list or the G-reg points for the initial global alignment is controlled by the "Point-based reg. type:" choice list; "LM based" specifies use of landmark list points and "G-reg based" specifies the use of the G-reg points. Obviously, if G-reg points are used, then they must have been marked for all database entries (i.e. identified in each database volume prior to clicking the "Landmarks->database" button), and so the user must consider the global registration technique during database preparation.

The software implements a number of tests on the point-based stage of global registration. These will ensure that at least four points are available and that they are sufficiently non-coplanar. In addition, the optimal alignment of any sets of points should bring the centroids of the point sets into alignment. This could be used to calculate the translation parameters of the transformation model directly, without optimisation. However, the software actually optimises all nine parameters of the transformation model. This allows the distance between the point set centroids after alignment to be used as an error estimate; it is displayed in the Automatch Database Inspector dialog box after global registration has been performed, and should in general be less than one voxel.

13.2.2 The G-reg Point Filler

The G-reg point filler provides functionality to copy locations from four points in the current landmark list to the four G-reg points. Enter the numbers of the landmark points to be copied into the "Greg 1", "Greg 2", "Greg 3" and "Greg 4" fields in the Automatic Landmark Point Placement tool, and then click the "Fill G-reg points" button. If any of the specified points cannot be copied e.g. a non-existent landmark point number is entered, then a warning will be issued.

The G-reg point filler can be used to accelerate the automatic landmark point placement process in two ways. First,

the user is required to manually identify all landmark points in all image volumes that will be entered into the automatch database. It may be more convenient to perform this process first, outputting the results to file for later use in database construction. The G-reg point filler allows the locations of any four of these points to be rapidly copied into the G-reg points during database construction. Second, once a database has been constructed it can be used to analyse an arbitrary number of further image volumes. However, the point-based stage of registration requires four landmark points to be identified in each of these volumes. It may prove more efficient to prepare a landmark point names file containing only four points, and to manually identify these points in each volume to be automatically analysed, saving the results to file. These files can then be loaded prior to automatic analysis of the corresponding image volume and the four point locations copied to the four G-reg points using the G-reg point filler. The landmark point names file for the full set of landmark points to be automatically located can then be loaded and the automatic landmark point location process applied.

13.3 Image-based Refinement of Global Registration

The point-based stage of global alignment provides the required initialisation for a second, image-based stage. This uses full-sized (i.e. spanning the full size of the volume from edge to edge) image patches intersecting at the centre point of the volume and aligned with its major axes. Simplex optimisation is used to optimise the parameters of a nine-parameter affine transformation through a maximum-likelihood method, based on the likelihood function described above. The γ factor is set to 1 during global registration; it is then calculated from the aligned, full-sized patches once the global registration is complete, and used during local registration (this procedure is based on the theory underlying the expectation-maximisation algorithm). The smoothing applied during the image-based stage of global registration is controlled by the "Glob. sig." field in the Automatic Landmark Point Placement tool; this specifies the standard deviation of the smoothing kernel in voxels, in the down-sampled coordinate system i.e. the coordinate system of the data currently loaded into the database and the Sequence tool. As mentioned above, the "Glob. border" field must have a value at least three times the size of the "Glob. sig." field.

The software will display a progress bar during image-based global alignment. In addition, the number of iterations of the simplex optimisation and the value of the cost function at the end of optimisation will be displayed in the top-level tinaTool window. Once transformation models have been calculated for all database entries, they will be stored in the automatch database, and can be inspected in the Automatch Database Inspector dialog box. The χ^2 per degree of freedom at the end of global alignment is also displayed.

13.4 Checking the Global Registration Result

The results from the global registration stage of the software form the initialisation for the coarse-to-fine local registration of patches of image data around each landmark point. If this initialisation is not accurate enough, the local registration will fail to find the correct locations for the landmark points, and the automatic landmark point location software will perform poorly. Failures in the global registration stage of the software can indicate several different problems:

- The points used in the initial, point-based stage of the global registration are not consistent i.e. different points have been selected for different entries in the database, or for the current volume;
- The points used in the initial, point-based stage of the global registration are poorly located e.g. errors have occurred during manual landmarking or the points themselves vary significantly between different image volumes;
- The points used in the initial, point-based stage of the global registration are poorly selected e.g. they are close to coplanar, or cover a small volume of the total sample;
- The shapes of the specimens in the automatch database vary too much from the shape of the specimen in the current volume to allow a good, image-based alignment.

Therefore, the user should check the global registration result for a representative sample of the image volumes to which the automatic landmark location software will be applied.

In order to check the global registration result (assuming that an image volume and automatch database have been loaded, and global registration has been performed):

- Start the TINA Imcalc tool by clicking on the "Imcalc" button in the top-level tinaTool window. Start a new TV tool, and install it on the Imcalc TV.
- Open the Automatch Database Inspector dialog box (see Section 17).
- Check the value in the "Centroid error:" field; it should be less than 1. If the value is greater than 1, then the point-based stage of global registration has failed; this usually indicates that the G-reg points defined for the currently loaded image volume do not correspond to those defined for the specimen in the database, or that errors have occurred during manual landmarking of these points.
- Check the value in the "Chisq:" field; this is a rough indication of the quality of the image-based stage of global registration, since the specimens in the current volume and the database will have different shapes, and so no accurate guide can be given. However, the value should be less than around 30, and a value of less than around 15 indicates a high-quality alignment.
- Select the "RX" option in the "Image patch:" choice list of the dialog, and then click "Push". A complex image will be pushed onto the Imcalc stack and displayed in the Imcalc TV tool as an anaglyph image. The real part of this image, displayed in green, is the full sized y-z slice through the centre of the image volume stored in the current automatch database entry i.e. the entry currently displayed in the dialog box. The imaginary part, displayed in red, is the corresponding slice through the current volume, obtained using the transformation model optimised during the global registration. The two images should be approximately aligned; they will not align exactly due to shape differences between the specimens, but should be in the same approximate orientation and of the same overall size. Any gross misalignment indicates that the shape differences between the specimens are too great to allow a good alignment i.e. the specimen stored in the database cannot be used to automatically identify landmark points in the current volume.
- Repeat the previous step for the "RY" and "RZ" image patches.
- Repeat the process for each entry in the database: if there are not at least 6-7 specimens in the database with a good alignment to the current image volume, then the automatic landmark point location software will perform poorly.

If persistent problems occur with global registration, consider constructing a database from specimens more similar in shape to those that will be automatically analysed, or using a different selection of points in the initial, point-based stage of the global registration.

13.5 Coarse-to-fine Local Registration

The global registration provides an approximate alignment of the current image volume to those volumes represented in the automatch database. This acts as an initialisation for a local registration of smaller image patches around each landmark point. The local registration must both search the cost function for the global minimum and then refine the estimate of the global minimum for maximum accuracy. These two objectives are, in some ways, diametrically opposed, since efficient searching requires large patches and significant smoothing, whereas maximum accuracy requires small patches, in order to improve the fit of the model to the data, and low smoothing. Therefore, the local registration uses a coarse-to-fine strategy involving three stages with decreasing patch and smoothing kernel sizes. The stopping criterion for the first two stages is less strict than for the third stage, since only an approximate solution is required to initialise the following stage.

The local registration optimises only the translation and scaling parameters of the nine-parameter transformation model i.e. does not optimise the rotation. This avoids degeneracy when aligning landmarks located on surfaces of constant curvature; for example, a sector taken from a circle can be aligned to any other sector of that circle if rotation is allowed, but will only match its original position if the transformation is limited to translation and scaling only.

The patch and smoothing kernel sizes can be controlled using the "Loc. patch size 1", "Loc. sig. 1", "Loc. patch size 2", "Loc. sig. 2", "Loc. patch size 3" and "Loc. sig 3." fields; the number refers to the registration stage. Local registration is performed by clicking on the "Automatch" button. The software will automatically add a border around the specified patch size equal to three times the size of the smoothing kernel; the local registration will be aborted if the stored patch size in the database is too small to allow this. These images patches will then be aligned for all points in all records in the database; this should take approximately one second per point,

but will take considerably longer if the TINA libraries and the Geometric Morphometrics toolkit have not been compiled with parallelisation enabled (see Section 2). A progress bar will be displayed, informing the user of the current database record and point being aligned. Once local registration has finished, the projected location of each point in the database i.e. the estimated position of that point in the current volume based on the global and local registration results, will be displayed in the Automatch Database Inspector dialog box, together with the χ^2 per degree of freedom calculated from the third-stage patch using both the global registration result only, and the combined global and local registration results. The χ^2 per degree of freedom for the global plus local registration should be close to 1 if the point has fitted correctly.

The registration will abort if the global registration has not been performed i.e. there are no stored global transformation models in the database. All stored transformation models will be deleted if a new image volume is loaded into the Sequence tool, since they will no longer be valid. However, if the automatch database is stored to file then the results of all stages of registration are stored in that file. This allows all registrations to be optimised and stored for a specific current volume, facilitating rapid testing of the final refinement and error detection stages of the software. However, it also permits the situation where the user saves a database containing global and local registration results for a specific current image volume, loads a different volume into the sequence tool, and then loads the database for the original current volume, ending up with a set of global and local registrations for the wrong current volume. The software has no way to detect this error, and so the user must be careful to avoid it.

13.6 Final Refinement of Landmark Estimates

Once the global and local registration stages have been completed, the database will contain multiple estimates of the position of each point in the landmark list: one for each entry in the database. The purpose of the final refinement stage is to combine these estimates into a single hypothesised location for the landmark, the coordinates of which can then be entered into the current landmark list. However, as with all optimisation-based machine vision systems, there is a chance that some of the estimated locations contained in the database will be fit failures i.e. will be in the wrong location. This can be caused by several different effects: shape differences can result in the structures around the landmark point in the current volume being a poor model for the structures in the database entry, similar structures in different locations within the images can generate local minima in the likelihood cost function that trap the optimisation, and local minima can also be generated by image noise. Therefore, the final refinement must take account of the fact that some of the estimated locations for each landmark may differ significantly from the correct location i.e. may be outliers.

The potential presence of outliers in the database militates against using any procedure based on a logical AND process e.g. taking the mean or median of the estimated locations to produce the final hypothesised location for each landmark. Furthermore, taking the single estimated location with the lowest estimated error would be undesirable, for two reasons. First, lacking a model of shape variation between patches in the database and the current volume, any error estimate (such as the χ^2 per degree of freedom at the end of the local registration) is itself unreliable, which could lead to errors in point selection. More importantly, picking a single point could lead to systematic error, through marking up all points in a given dataset using the single database entry with the most similar shape, thus transferring any errors on the point locations in the database entry directly onto the dataset being marked up. This could exaggerate any shape differences observed in subsequent Procrustes analysis.

Instead, a procedure is required that uses a logical OR process i.e. finds the best single estimate based on the consensus of the estimates contained in the database. The final refinement achieves this by constructing a 3D space large enough to hold all of the estimated locations for a given landmark point, entering the points as Kronecker delta functions, and then convolving the space with a Gaussian smoothing kernel. The size of the smoothing kernel is set in the "Fin. sig." field (see Section 18), and is given in the coordinate system of the current sequence i.e. taking into account any down-sampling applied during data loading. The final refinement algorithm selects the location of highest point in the space after the convolution i.e. the most probable point location based on any sub-set of the estimated locations from the database, as the final estimated point location. If no highest point exists, since the estimated locations are all separated by more than three times the size of the Gaussian kernel used in the convolution, then the point from the database with the lowest χ^2 is used instead and a warning message is displayed in the top-level tinaTool window. This should be considered an error condition for the reasons outlined above.

The "Fin. sig." field should be set to a value representing the expected accuracy of the registration result for points that successfully find the correct location i.e. ignoring any outliers). One way to determine this value is to manually identify landmarks in a representative image volume twice; the median absolute error in voxels between the manually identified positions for each landmark provides a guide to the achievable landmarking accuracy, and can be used as the value for "Fin. sig.". However, the user must remember that the landmark positions output

to file are in the original coordinate system of the data i.e. ignoring any down-sampling applied during loading, whilst the "Fin. sig." parameter is in the coordinate system of the current sequence i.e. takes any down-sampling applied during loading into account.

The final refinement has two additional options. First, the "Chi-sq. weight" option weights the Kronecker delta functions using the χ^2 at the end of the local registration i.e. down-weights the points that may not have fitted correctly. This option requires careful testing before use, since the estimated χ^2 values themselves are not completely reliable, and so it is recommended that "Off" should be selected. Second, the "Snap points" option will take the location predicted by the final refinement, and then move it to the closest surface as defined by the "Bone threshold" option in the Manual Landmark tool.

The final refinement can be run by clicking the "Gauss conv." button. A progress bar will be displayed; however, for small Gaussian kernels the process will be completed relatively quickly. The point locations predicted by the final refinement will be copied into the current landmark list, replacing any coordinates already stored there. A current landmark list must exist in the software i.e. at least a landmark point names file must have been loaded. The type code of each point will be changed to 2, so that the user can easily see which points in the landmark list have been automatically estimated.

13.7 Outlier Detection

With careful database construction, the automatic point location software is capable of achieving point location errors equivalent to those from manual point identification for 95% of the points in typical landmarking experiments. However, occasional outliers will remain, and so the user is expected to manually inspect automatic point locations and refine them where necessary (this should be much faster than manually identifying all points). The outlier detection stage of the software provides a guide to this process, attempting to identify these outliers and indicate them to the user.

Lacking a model of shape variation between the image volumes in the database and the current volume, no stage of the automatic point location can be completely accurate. Therefore, errors may exist on both the automatic point locations and the estimated errors on those locations provided by the χ^2 per degree of freedom. The outlier detection therefore inspects consistency across both estimated location and estimated error, so that an error on either quantity will result in a point being flagged as an outlier. A leave-one-out experiment involving 14 micro-CT datasets of mice, wood mice, voles, hamsters and gerbils indicated that this provides a reliable outlier detection method with a false negative rate of 0%, at the expense of a false positive rate of 20% on image volumes for which the majority of the database consists of examples from the same genus.

The outlier detection operates by sorting the predicted locations for a given landmark point from the database into order, based on their χ^2 per degree of freedom at the end of local registration, and then flagging any point as an outlier if the estimated locations for the points with the lowest χ^2 differ from the location resulting from the final refinement stage by more than a given threshold. The number of points included in the check is controlled by the "Error checks:" check list; "P1" checks the point with the lowest χ^2, "P2" the point with the second lowest χ^2, and "P3" the point with the third lowest χ^2, The distance thresholds are set in the "P1 thresh.", "P2 thresh." and "P3 thresh." fields. Click the "Error analysis" button to run the outlier detection. If a point fails outlier detection, an error code will be added to its type code based on translating a binary mask of failure flags from each stage into an integer e.g. failure of stages 1 and 2 results in the binary mask 011, which has integer value 3, resulting in a final type code of 5 (since the type code will be 2 after automatic point location). Points with type codes indicating an failed check will be displayed in the 3D TV as green and red chequered cross-hairs or spheres. The user should check the location of each of these points.

Note that, unlike the other parameters in the Automatic Landmark Point Placement tool, the error thresholds are given in the original coordinate system of the data, not the down-sampled coordinate system of the current sequence. There is a reasonable direct proportionality between the error thresholds and the maximum error on landmark points that will pass error checks at those thresholds. Therefore, giving the error thresholds in the original coordinate system of the data allows the user to set them to the maximum acceptable error on landmark point locations in the same coordinate system used in the output files. As a rough guide, the error thresholds should be set to approximately the same value as the patch size parameter used in the last stage of local registration, taking this coordinate system transformation into account, in order to avoid false negatives. For example, if the "Loc. patch size 3" parameter is set to 10 but the sequence was loaded with "Stride" and "Downsample" set to 2, then the error detection thresholds should be set to around 20.

The outlier detection requires that all previous stages of the automatic point location have been applied, and will not operate if any manual alteration of the automatically located points has been performed. This is enforced by

checking the type code of all points in the current landmark list; they must all be of type 2. The user can set these type codes through the Landmark Points dialog box to work around this check, but does so at their own risk.

13.8 Correction of Outliers

The parameters of the outlier detector should be tuned to provide a minimum number of false positives i.e. points that, whilst in the wrong location, are not flagged as outliers. This reduces the probability that erroneous points will be passed to subsequent Procrustes analysis, but inevitably means that a large number of false negatives will be generated i.e. points that, whilst in the correct location, are flagged as outliers. However, such points require only brief inspection by the user and an alteration of their type code to indicate that they are in the correct position. This can be done efficiently by making use of the "Jump lock:" option in the Manual Landmark tool.

A complete set of steps required to efficiently correct outliers in the list of automatically located landmark points is as follows:

- Switch the "Jump lock:" option in the Manual Landmark tool to "On".
- Scroll to the start of the landmark list by clicking the "Start" button in the Manual Landmark tool or the Landmark Points dialog box.
- Click the "Next outlier" button. The current landmark will be advanced to the first landmark in the list that has been identified as a potential outlier, indicated by a type code of 3 or greater. Since the "Jump lock:" option is on, the 3D cursor will be moved to the automatically identified position for this landmark, and all TVs will be updated accordingly.
- Inspect the indicated position of the landmark. If an error has occurred, move the 3D cursor to the correct position using the arrow keys in the Manual Landmark tool, or by using mouse or keyboard interaction with the images in the TVs.
- Click the "Mark Point" button in the Manual Landmark tool, or use the keyboard shortcut assigned to this button in any TV. The (corrected) coordinates of the 3D cursor will be transferred to the landmark point and the point type code will be changed to 1 to indicate that the point has been manually identified. Note that, even if the point was in the correct position and has not been adjusted, clicking on the "Mark Point" button in the Manual Landmark tool, or using the keyboard shortcut assigned to this button in any TV, will update the type code to 1 to indicate that user inspection has occurred and the point is in the correct location.
- Proceed through the landmark list by clicking the "Next outlier" button, inspecting points, correcting them if necessary, and clicking the "Mark Point" button whether or not any correction has been applied, until the end of the landmark list has been reached.

14 Using the Automatic Landmarking Tool

A complete set of steps required to automatically identify a set of landmarks is as follows:

- Prepare the database as described in 13.1, by manually identifying a specific set of landmarks in multiple specimens of similar shape to the image volume intended for automatic landmarking and, if desired, the G-reg points. Click the "Landmarks->database" button once all points have been identified in each volume, before proceeding to the next volume.
- Load the image volume intended for automatic landmarking.
- Load the landmark list names file used to identify landmarks for the volumes in the database.
- Manually identify at least four points in the landmark list or, alternatively, the four G-reg points.
- Ensure that the "Point-based reg type:" choice list is set to the correct option, depending on whether G-reg or landmark points were identified in the previous step.
- Ensure that both "Point based" and "Image based" are selected in the "Reg stages:" check list.
- Click the "Global reg" button to perform global registration.
  - Check the global registration result using the procedure described in Section 13.4.
- Click the "Automatch" button to perform local registration.
- Click the "Gauss conv." button to perform final refinement.
- Click the "Error analysis" button to run the outlier detection.
- Inspect the positions of any points that have failed outlier detection, correcting them where necessary, using the procedure described in Section 13.8.

14.1 Setting the Free Parameters

The automatic landmark point placement software depends on multiple free parameters: the patch sizes for the local registration, the smoothing kernel sizes for both the local and global registrations, the convolution kernel size for the final refinement stage, and the thresholds used in outlier detection. The default values were optimised for micro-CT images of rodent skulls with a resolution of 658×658 voxels within each slice, and containing 1000 to 1200 slices, loaded with the "Stride" and "Downsample" factors set to 2. If the software is applied to data that varies significantly from this description, then the user may have to optimise the free parameters in order to ensure proper operation of the software. This can be done using a leave-one-out experiment, as follows. Obtain 10 to 15 representative image volumes, and manually identify the desired set of landmarks in all of them. Then run the automatic point location software for each image volume, using the others to construct the database. Analyse the results, comparing the automatically estimated landmark point locations for each volume to the manually identified landmark point locations in order to estimate the errors. Repeat the process, varying the free parameters, in order to determine the optimal parameters for the data. This procedure involves large numbers of experiments, and so will be considerably faster if the TINA macro file system is used to run the experiments (see Section 15). Furthermore, the outlier detector operates by placing thresholds on the distances between the landmark point location determined by the final refinement stage of the algorithm and the locations of the three corresponding points in the automatch database with the lowest χ^2 per degree of freedom after global and local registration. Since all of this information is contained in the point positions files (see Section 17.1), the optimal error detection thresholds can be determined by outputting such files and using them to construct ROC curves.

15 "Double Pass" Operation of the Software

As described in Section 13.2.1, the global registration stage of the automatic landmark point identification software incorporates an initial stage based on user-identified landmark points, to deal with any gross rotation of the specimen relative to the previous specimens entered into the automatch database. This introduces a requirement for user interaction at the start of the automatic landmarking process. However, in large landmark-based geometric morphometric studies, involving hundreds of specimens, it may be desirable to eliminate this stage of the analysis, so that no user interaction with the software is required prior to checking and correcting any outliers. "Double pass" operation of the software allows this, but relies on careful positioning of the specimen in the CT scanner, such that:

- All specimens involved in the analysis, including those used to build the automatch database, are in approximately the same location and orientation in the image volumes;
- All specimens are of approximately the same size;
- Each image volume contains only one specimen.

Double pass operation of the software involves running the automatic point location procedure twice. The first pass is performed with no initial, point-based registration. The number of automatically located points that pass error checks in this process will be relatively low (typically 60% of the points in tests with relatively well-aligned micro-CT images of mouse skulls). However, these points can be used to perform the point-based stage of registration in a second pass through the automatic point location process, resulting in a more accurate global alignment and thus a higher proportion of automatically located points that pass error checks (typically 80-90% of the points in tests with relatively well-aligned micro-CT images of mouse skulls).

The double pass operation of the automatic point location software:

- Load the data set in which landmarks are to be automatically located, the automatch database to be used, and the landmark points names file.
- Deselect the "Point based" option in the "Reg stages:" check list of the Automatic Landmark Point Placement tool, and select the "Image based" option.
- Run the automatic landmark location software by clicking on "Global Reg", "Automatch", "Gauss conv." and "Error analysis" in turn in the Automatic Landmark Point Placement tool.
- Select both the "Point based" and "Image based" options in the "Reg stages:" check list of the Automatic Landmark Point Placement tool.
- Select "LM based" in the "Point-based reg. type" choice list of the Automatic Landmark Point Placement tool.
- Run the automatic landmark location software a second time by clicking on "Global Reg", "Automatch", "Gauss conv." and "Error analysis" in turn in the Automatic Landmark Point Placement tool.
- Check and correct any remaining outliers using the procedure given in Section 13.8.

Prior to using the software in this way, users should check that the specimens in their scans are well enough aligned by constructing an automatch database, then loading a few example scans and running the global registration stage of the software, checking the results as described in Section 13.4.

TINA incorporates a macro system that allows the user to construct a text file listing TINA operations. This system can be used to construct multiple macro files, each of which performs the analysis listed above on a single image volume and then outputs the results in Raw file format. An arbitrary number of samples can then be analysed e.g. overnight by running the macro files as a batch job; the only user interaction required is to load each image volume and Raw format landmarks file once the automatic analysis is complete, and check or correct any points identified as potential outliers. See the TINA Users Manual [19] for more details on the construction of macro files.

16 Approximate Projection

It may be useful, under some circumstances, to have a set of very approximate point locations that, whilst not accurate enough for shape analysis, can instead be used to initialise a manual landmark procedure. The software provides this option through the "LM projection" button. This will take the set of landmark points in the database corresponding to each point in the landmark list, project them into the current volume based on any registration results available (i.e. point-based global, image based global, and local; whichever stages have been performed), and find the mean of these projected positions, writing the result to the current landmark list. The "Projection:" check list provides two options for this procedure. "All " will replace all of the points in the landmark list, even if some of them have been manually identified; un-select this option to avoid overwriting landmark points that were manually located for the point-based stage of the global registration. "Snap" will take the mean projected location and then move it to the closest surface defined by the "Bone threshold" field in the Manual Landmark tool before entering it into the current landmark list.

The landmark projection can operate using only global registration results, making it much faster than the full automatic point location procedure. However, it is not considered to be automatic point location, since all results will have to be manually refined, and so the point type code is not changed by this procedure.

17 Automatch Database Inspector

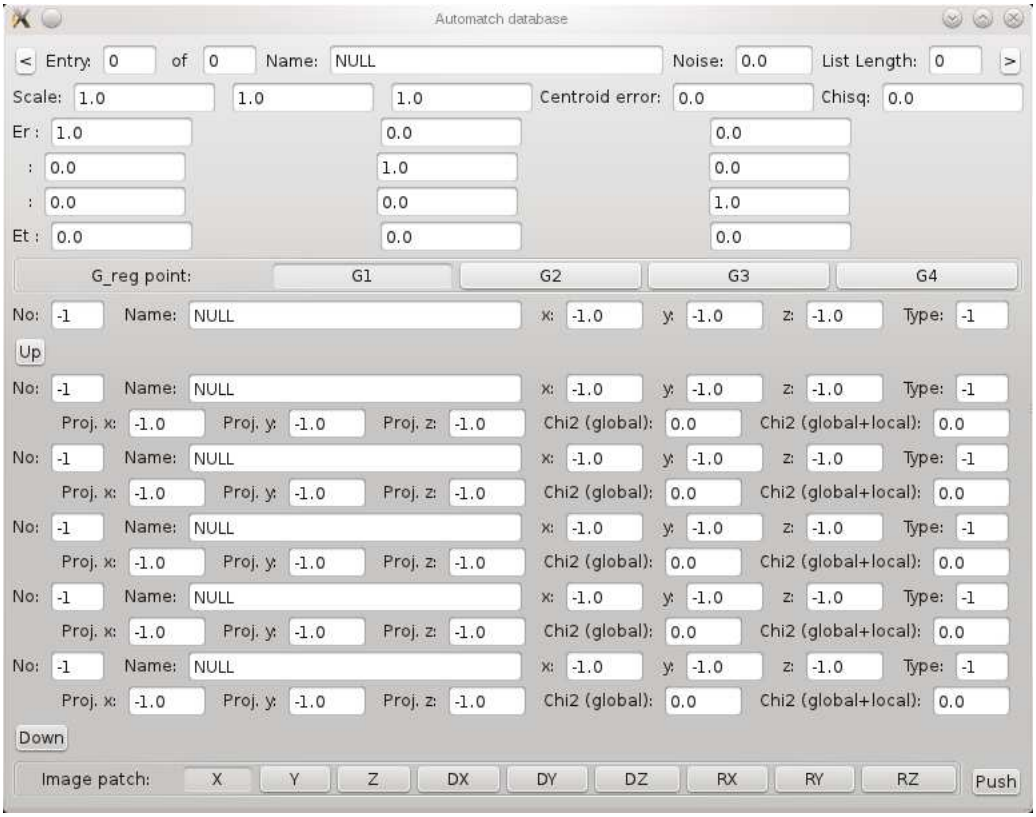

Figure 18: The Automatch Database Inspector dialog box.

The Automatch Database Inspector dialog box allows inspection of the information stored in the automatch database. It can be started by clicking on the "Automatch DB Inspector" button in the Automatic Landmark Point Placement tool, and is shown in Fig. 18.

The < and > keys allow the user to scroll through the entries in the database; the total number of entries and the ordinal number of the one currently being displayed are indicated in the "Entry" field. The "Name" field gives the filename of the image volume used to create the entry. The "Noise" field displays the estimated standard deviation of the image noise; this is used in the calculation of the likelihood function for image-based registration. If global registration has been performed, the "Et", "Er" and "Scale" fields display the parameters of the optimised transformation model (the rotation is displayed as a rotation matrix although, in effect, only three parameters are optimised; the optimisation itself uses a quaternion representation of the rotation parameters in order to avoid problems with gimbal lock). The "Centroid error" field provides an estimate of the error on the point-based stage of global registration, using the distance in voxels in the coordinate system of the current volume (i.e. including the effects of any down-sampling applied during data loading) between the centroids of the registration point sets in each volume. The "Chisq" field gives the χ^2 per degree of freedom at the end of the image-based stage of global registration, which provides an estimate of how well the transformed, full-sized patches through the centre of the current volume model the equivalent patches stored in the database entry.

The "G-reg point" choice list allows information about the G-reg points to be displayed, if any are stored. The options "G1", "G2", "G3" and "G4" select which point will be displayed, and the "No", "Name", "x", "y", "z" and "Type" fields in the next row of the tool display the point number, name, coordinates and type code.

The next section of the dialog displays information about the individual landmark points contained in the database entry. The "List Length" field displays the number of points. As in the Landmark Points dialog box, information about five points is displayed simultaneously, and the user can scroll through the list of points using the "Up" and "Down" buttons. For each point, the point number, name, and type code are displayed. The "x", "y" and "z" fields display the coordinates of the point in its original image volume; these are in the original coordinate system of the data i.e. ignore any down-sampling that was applied during data loading. If any registration has been performed, then the "Proj. x", "Proj. y" and "Proj. z" fields display the coordinates of the point projected

through the optimised transformation model i.e. the estimated coordinates of the corresponding point in the current volume. The "Chi2 (global+local)" field displays χ^2 per degree-of-freedom at the end of local registration; the "Chi2 (global)" field displays the equivalent (i.e. calculated from a patch of the same size as the one used in the last stage of global registration) but using only the global registration parameters. Any error score that has not yet been calculated has a default value of zero.

The last section of the dialog allows the user to inspect any image patches stored in the database. The "Push" button will push the selected image onto the Imcalc stack, allowing it to be viewed in the Imcalc TVs. "X", "Y" and "Z" select the local patches for the current point of the current database entry (the current database entry is the one being displayed; the current point is the one displayed on the middle two rows of the ten rows of landmark point information), in the form in which they are stored in the database. "DX", "DY" and "DZ" select the local patches for the current point of the current database entry in the form in which they are used in the image-based registration i.e. with smoothing and differentiation applied. This generates two images, the derivatives with respect to the x and y directions of the image patch; the y-derivative image is pushed onto the stack first, and so the x-derivative image will be on the top of the stack and the y-derivative image will be second on stack once the process has been completed.

"RX", "RY" and "RZ" select the full-sized patches used in global registration. If an image volume is loaded into the sequence tool then the corresponding patch from the current volume will be extracted and combined with the patch to form a complex image, in which the real part is the patch stored in the automatch database and the imaginary part is the corresponding patch extracted from the current volume. If global registration has been performed, then the result will be used when extracting the corresponding patches from the current volume. The complex image is then pushed to the Imcalc stack, and so can be viewed as an anaglyph image in the Imcalc TV. The real part of this image, displayed in green, is the full sized y-z slice through the centre of the image volume stored in the current automatch database entry i.e. the entry currently displayed in the dialog box. The imaginary part, displayed in red, is the corresponding slice through the current volume, obtained using the transformation model optimised during the global registration.

In all cases, X, Y and Z refer to the axis of the original volume normal to the plane of the image patch e.g. the X image is generated on a y-z plane in the image volume.

17.1 The Database Positions Output File

During automatic landmark point location, projected coordinates for all landmark points in the automatch database are stored after each stage of registration. The Automatic Landmark Point Placement tool supports the output of files containing any of the sets of coordinates currently stored in the automatch database, in ASCII text format. Such files can be useful in analysing the operation of the software and in setting some parameters, such as the error thresholds for the outlier detection stage.

The set of coordinates that will be output to file is controlled by the "Pos. output field:" choice list in the Automatic Landmark Point Placement tool.

- Proj. glob.: the coordinates of the point after transformation using the result of the global registration.
- Proj. 1: the coordinates of the point after transformation using the result of the global registration and the first stage of the local registration.
- Proj. 2: the coordinates of the point after transformation using the result of the global registration and the first and second stages of the local registration.
- Proj.: the coordinates of the point after transformation using the result of the global registration and all three stages of the local registration.
- Pos: the coordinates of the point without any transformation i.e. the manually identified location of the point in its original image volume.

Clicking on the "Pos. output" button will write these locations to the file specified in the "Output:" field of the Automatic Landmark Point Placement tool.

Each line of the point positions file corresponds to one entry in the automatch database. The first two numbers on the line are the centroid error and χ^2 per degree of freedom for the global registration. The remainder of the line contains sets of six numbers corresponding to each landmark point in the database entry. The first three of these are the x, y and z coordinates of the point; the set of coordinates that will be used is controlled by the "Pos. output

field:" choice list as described above. The fourth number is the χ^2 per degree of freedom of the match between the stored local image patches around the point, using the patch size specified in the "Loc. patch size 3" field, and the corresponding patches around the transformed point in the current volume using the transformation model from the global registration i.e. prior to local registration. The fifth number is the χ^2 per degree of freedom of the match between the stored local image patches around the point, using the patch size specified in the "Loc. patch size 3" field, and the corresponding patches around the transformed point in the current volume using the full transformation model including both the global and local registration results. The sixth number is the point type code. The final line of the file contains equivalent information from the current landmark point list. For example, if the final refinement stage of the automatic point location algorithm has been performed, then the automatically estimated point locations will be written to the file; if the outlier detection stage of the automatic point location algorithm has been performed, then the point type codes on the last line of the file will indicate any failed error checks.

In order to simplify the operation of loading the point positions files into analysis packages such as Mathematica, all lines in the file are of the same length, and any missing values are replaced with zeros. For example, there is no centroid error or χ^2 for the current volume; therefore, the first two numbers on the last line of the file are always zero. Therefore, the number of rows in the file is given by

$$rows = N + 1$$

where N is the number of entries in the database, and the number of columns is given by

$$columns = 6n + 2$$

where n is the number of points in the landmark list.

18 Notes on Future Updates

This section lists some ideas for future updates to the software that could improve usability with little effort, but which have not been implemented due to time constraints.

18.1 Initial Global Alignment

The point-based global alignment is a placeholder used during software development, since it allows rapid testing of the software when automatically identifying landmarks for comparison to manually identified, gold-standard landmark locations. However, it is arguably not the ideal method for user interaction with the software.

Fortunately, it would be possible to implement a much more convenient method. The user could be prompted to rotate the view in the 3D TV to a fixed orientation; it does not matter what this orientation is, as long as it is repeatable across different specimens to reasonable accuracy. The viewing transform could then be obtained (call get-viewing-transform() in tlmedMousetool-volpack-interface.c) and stored in the automatch database. The combination of the viewing transform from the database with one from the current volume would provide the rotation parameters for an affine registration of the two; it would then be trivial to obtain the translation and scaling in an image-based registration limited to only those parameters. Alternatively, the bone threshold in the Manual Landmark tool could be used to identify a set of surface points, the centroid and standard deviation of which would provide the translation and scaling parameters. A second image-based stage of global registration could then be performed to refine the transformation model. This procedure would avoid the need to mark up any landmark points prior to automatic point location.

18.2 Covariance Estimation

The covariances of the parameters optimised in any maximum likelihood technique can be estimated using the minimum variance bound, as long as the likelihood is properly normalised. "Properly normalised" can be defined using

$$-2lnL = \chi^2 \tag{3}$$

i.e. normalised such that the distribution of $-2lnL$ is given by the χ^2 distribution, [4, 2, 3, 5]. Therefore, the maximum likelihood optimisation can simultaneously provide optimised model parameters, an estimate of the goodness-of-fit of the model in the form of the χ^2 per degree-of-freedom at the end of optimisation, and an estimate of the errors on the optimised parameters in the form of the minimum variance bound. Of course, the errors given by this technique represent the extent to which the estimated parameters are destabilised by the noise on the data, and will only correspond to the total error on the parameters if the model is a good fit to the data (i.e. the residuals consist only of noise) and the optimisation procedure has successfully located the global minimum of the cost function.

A parameter covariance calculation function based on the minimum variance bound is included in the software (compute-covar() in tlmedMousetool-patchmatch.c). However, calculation of the minimum variance bound can be problematic in practice; for example, in [2, 3, 5] it was found necessary to calculate the derivatives of the likelihood function over a range of values, and then take the median, in order to get a stable estimate. The function currently included in the software has been subjected to basic testing (i.e. it runs and produces the correct single-point estimates of the minimum variance bound). However, more rigorous stability testing and, possibly, modification in order to improve numerical stability, is required prior to use.

The parameter covariances could be used in two ways. First, they could be included in the output files and thus made available for use in later analysis. However, they could also be used to set the size of the Gaussian kernel used in the final refinement stage of automatic landmark location; this would provide a more optimal combination of the available data and thus a more accurate final hypothesised landmark location.

18.3 Bone Threshold Estimation

Several functions in the software depend on the value entered into the "Bone threshold" field in the Manual Landmark tool, including mouse interaction with the volume rendering displayed in the 3D TV in "Pick" mode, landmark point display in the 3D TV, and the various "Snap" functions in the Manual Landmark tool and the Automatic Landmark Point Placement tool. The procedure required to set this value is given in Section 5.2.2. However, using a single value is sub-optimal for several reasons. For example, the true threshold for points on bone surfaces in typical geometric morphometrics studies varies across the image, and is different e.g. at the interface

between tooth enamel and air space compared to the interface between soft tissue and bone. The software has been designed to avoid consequent problems e.g. by encouraging the user to refine manual landmark points using the 2D TVs, in order to ensure that the use of a simple threshold does not result in any bias on the landmark locations.

TINA includes an advanced medical image segmentation algorithm [6, 7, 8, 1] that could be used, in the first instance, to set the bone threshold automatically. However, it could be used in a more advanced fashion, to analyse the intensity profile along any vector through the image data during landmark point selection using the 3D TV in "Pick" mode and automatically locate interfaces between any pair of tissues, removing the need for a simple intensity threshold and avoiding any consequent bias on landmark point locations. Finally, the optimised tissue model produced by the TINA medical image segmentation algorithm could be used to enable the multi-material rendering options in the Volpack library allowing, for example, soft tissues to be displayed as a transparent surface, in a different colour, above the bone surface in the 3D rendering. This would allow the user to consider biomechanical aspects of the specimen, such as the size of particular muscle groups.

Part 3: The Shape Analysis Tool

H. Ragheb and N.A. Thacker
Imaging Science and Biomedical Engineering Division,
Medical School, University of Manchester,
Stopford Building, Oxford Road,
Manchester, M13 9PT.

19 Introduction

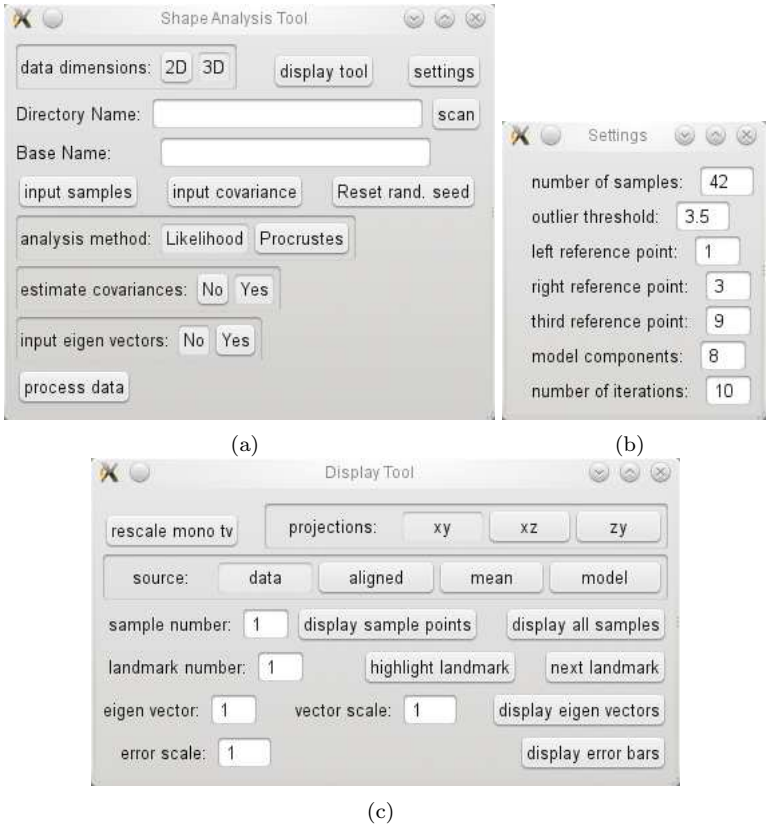

Figure 19: The Shape Analysis tool (a), the Settings dialog box (b) and the display dialog box (c).

The purpose of the Shape Analysis tool, shown in Fig.19, is to utilise the analysis of shape data with anisotropic measurement covariances, as a tool to monitor manual and automatic placement of landmarks. These methods are available via the automatic 3D landmarking tool, as a system for quality assessment and validation of output data.

19.1 Data

The standard input data file format for this software is TPS. Note that the extension must only consist of capital letters. For the case of 2D/3D data, the standard TPS file provides two/three coordinates for each landmark point (LM/LM3). However, one can provide the software with 3D data and force it to treat the 3D data as 2D, by selecting 2D as 'data dimensions'. In this case, the third coordinate values are ignored. When generating outputs of the TPS file format, the software writes two or three coordinates according to the current number of 'data dimensions'.

19.2 Methods

The methodology behind the software functionalities have been described in detail in [17] and [15]. Although in our technical report [17] it is assumed that data has two dimensions (x and y), the methods are applicable to 3D data with minor changes. The extension of our methods from 2D to 3D has been explained in [16].

19.3 General Use

The shape analysis tool can be used to obtain a number of outputs, results and conclusions as follows. These have been explained in details with discussions on the results obtained from example data sets in our reports [15] and [16]. We shall provide the list of steps needed to perform a number of experiments (Boxes A, B and C) later on.

1. Shape alignment by constructing a linear model and fitting the data based on an iterative optimisation approach.
2. Estimating covariances on each landmark based on computing sample covariances and applying covariance corrections.
3. Generating isotropic (whitened) output data for use later as input to popular packages which support a Procrustes analysis.
4. Comparing the covariance estimates with the ones from reproducibility studies based on repeat manual mark-ups.
5. Using the anisotropic measurement covariance output data as input data to process independent repeat data, where χ^2 ratios are used to assess statistical equivalence.
6. Using the list of potential outliers detected in the data to identify problematic data samples, for subsequent elimination or manual correction.
7. Processing Monte-Carlo data using the model (the mean and eigenvectors) together with the covariance estimates used when generating them, to generate plots of expected versus obtained measurement errors. This is the less challenging Monte-Carlo test.
8. The Monte-Carlo data can also be processed by constructing an independent model and estimating covariances accordingly. This is the more challenging Monte-Carlo test.

20 Software Guide

In what follows we outline how to use different dialog boxes and buttons to perform shape analysis experiments.

20.1 Shape Analysis Tool

In this section we describe how to work with the buttons appearing on the Shape Analysis Tool dialog box.

'data dimensions': One must select the dimensionality of the input data before pressing any other buttons. This tells the software whether each landmark is '2D' (x and y components only) or '3D' (x, y and z components). The default dimension is 3D. If 3D is selected but a 2D data file is provided, data will not be accepted. However, if 2D is selected while a 3D data file is provided the first and second coordinates of each point will be used and the third coordinate will be ignored.

'input samples': A directory needs to be created for every specific experiment on existing or new data. The corresponding TPS file should be copied into the directory. This leads the software to define this as its working directory (read/write). The file and its directory (full path) is located using the 'scan' button which provides a browsing facility. Appropriate messages will appear on the Tina tool window informing user about the success or any possible problems, should they occur. For instance if there are missing coordinates in the data, the corresponding samples will be removed from the data and a warning message confirms the action.

'input covariances': If there is a file containing covariance data in the format expected and consistent with the data dimensions and number of landmarks, it is possible to locate this file via the browser and input these covariance data and use them rather than estimating new ones. Although it is possible to generate these files manually, it is strongly recommended that no other covariance file is used other than those generated by the software for similar shape data (identical number of landmarks and dimensions).

'Likelihood/Procrustes': The default method is 'Likelihood' which is actually the method for which we have developed the software. This generates statistically valid estimates of parameters in a way which applies the Likelihood approach in a self-consistent manner (see technical reports). Alternatively, one can select the traditional 'Procrustes' method for which no covariance data will be estimated/used, this has been included largely for compatibility with established methods within the morphometrics literature. Output files from the Likelihood analysis are written (see below) which transform the data suitable for analysis in conventional Procrustes packages, by enforcing the requirement of homogenous measurement noise. This data set is already aligned, so that subsequent Procrustes alignment should not modify relative sample positions, but the new data now have homogenous error characteristics, suitable for use in PCA.

'estimate covariances': No matter whether the software has received covariance data or not, the user can tick 'Yes' to estimate covariances (default) or 'No' to use the covariance data provided already as input. In the latter case, if no covariance file has been provided, the identity matrix is used as the default covariance matrix. In the case of using Procrustes, even if a covariance file is provided, it will not be used.

'input eigenvectors': Here 'Yes' selection is specifically designed for one form of Monte-Carlo test where the Monte-Carlo data are fitted to the existing model. So, this experiment can only be done if the original data have already been processed. The files corresponding to the model (the mean and eigenvectors) are automatically generated by the Likelihood analysis and made available in the data directory every time a data set is processed by estimating the covariances. Hence, for this Monte-Carlo test, the covariance file should also be provided. The default selection is 'No'.

'rest rand seed': This button only has an effect on the Monte-Carlo data generation routine. By pressing this button, one can reset the seed value (using the system time) for the random number generator. Once the seed value is changed the software can generate a new set of Monte-Carlo data which is different from a previous set (before resetting the seed).

'settings': Before pressing the 'process data' button, the user must set all processing choices (Settings dialog box) appropriately. These are clearly explained below.

'process data': By pressing this button the process starts based on the choices made using other buttons of the tool.

'display tool': This tool box provides the means needed to display data and results graphically. Detailed descriptions are given below.

20.2 Settings

Here we explain what is the purpose of different settings appearing on the Settings dialog box.

'number of samples': This can be set to a number smaller or equal to the number of samples which have been successfully read from the TPS file. This is useful if user decides to experiment with a small subset of samples rather than the whole population. Note that, for the analysis to be statistically valid, depending on the data a minimum number of samples is needed (e.g. 30 samples).

'outlier threshold': This is the standard deviation value used as threshold in the outlier detection routine. This is usually set to a value between 3.0 and 5.0 depending on the data under study.

'left reference point': Enter here the left landmark point of the base line. Note that each landmark number is between one and the number of landmarks.

'right reference point': Enter here the right landmark point of the base line.

'third reference point': This landmark point is not needed for the case of 2D data, but must be entered appropriately for the case of 3D data to form a planar reference together with the left and right points.

'model components': This is the number of eigenvectors/values used when constructing the linear model. Obviously, one needs to have a good idea of the optimum number of model components for each data set, as using too many components would result in over-fitting and using insufficient number of components would result in the model failing to fit well to the data. This number is usually smaller for the Likelihood method than that for the Procrustes. We outline an experiment in Section 2.4 and Box C that helps users to find an optimum number for each data set as the number of model components.

'number of iterations': From our experience, 10 iterations (default) is sufficient to reach convergence. However, this number may be set based on an idea of how light/heavy the data set can be. Note that a minimum of 2 iterations is needed for the methods to converge.

20.3 Display Tool

In this section we explain how to use different buttons appearing on the Display Tool dialog box.

'rescale mono tv': To use the display tool a display window should be prepared first by pressing the Mono button and then the New Tvtool button on the Tina dialog box, followed by pressing the Install button on the New Tvtool. Finally press the 'rescale mono tv' button on the display tool to rescale the data in order to improve visualisation and its display coverage. Also it is recommended to select the largest size available (768) on the display (mono) window. Note that, when there are negative coordinates in the 3D mark-up data (e.g. Monte-Carlo data or data

generated by Morphoj software), one may experience difficulties to rescale the mono tv properly and be able to view all parts of the projections on the 2D planes. In such cases, one should press 'Mouse' button on the 'mono' tv and select 'zoom', then adjust the position and scale of the tv display in the usual way.

'projections': Selecting each of the projection planes ('xy', 'xz' and 'zy') determines which plane is currently active, so that the display buttons would only work on the selected plane. This is mainly for the case of 3D data. When working with 2D data, selection of planes does not matter and the 'xy' plane (default) is always active.

'source': Select which source of point data you want to be used in the display routines. Select 'data' for the input data, 'aligned' for the final aligned data, 'mean' for the final mean shape, and, 'model' for the final linear model constructed.

'sample number': The user may enter the sample number to display individual samples (between one and the number of samples).

'display sample points': By pressing this button, the landmark points of the sample number chosen will be displayed.

'display all samples': By pressing this button, the landmark points of all the samples studied are displayed superimposed. This is a useful button when used to display the input data before starting the process or to display aligned data and models once the process is over. Looking at such data spreads gives an idea of the quality of alignment.

'landmark number': The user may enter a landmark number to be highlighted on the display window (between one and the number of landmarks).

'highlight landmark': By pressing this button, the landmark point of the sample number chosen will be highlighted.

'next landmark': By pressing this button, the next landmark point of the sample number chosen will be highlighted.

'eigenvector': The user may enter the eigenvector number to be displayed individually (between one and the number of model components).

'vector scale': This is the scale value used to magnify the eigenvectors in order to improve visualisation on the display window, as the absolute values of vectors may be too small to see.

'display eigenvectors': By pressing this button, the eigenvectors are displayed taking into account the vector scale. It is recommended that the source is selected to be the 'mean' shape. One may display eigenvectors corresponding to several model components superimposed by changing the 'eigenvector' number and pressing 'display eigenvectors' repeatedly.

'error scale': This is the scale value used to magnify the error bars in order to improve visualisation on the display window, as the absolute values of errors may be too small to see.

'display error bars': By pressing this button, the error bars are displayed taking into account the error scale. It is recommended that the source is selected to be the 'mean' shape. One useful graph may also be obtained by displaying the aligned data superimposed. In the case of Procrustes, errors are only computed after the alignment and using the residuals left, which can be displayed in a similar manner.

20.4 Input/Output Files

It is strongly recommended that for every new experiment a new directory is created and the input data file is copied there. This is to make sure all the input and output files are the correct ones for the specific settings provided for the specific experiment. The files CovarianceFile, WhiteData, MonteCarlo, MeanShape, EigenVectors, EigenValues, ChisqAdjustFile and ChisqTest only correspond to the Likelihood analysis and are not generated/used when working with Procrustes.

TPS: This is a common file format in geometric morphometrics for reading and writing mark-up data, and our software expects this format when reading samples from the input file (while the file name is arbitrary). The file extension must be in capital letters (TPS).

SampleLabels: This file contains the labels (IDs) corresponding to the input samples received by the software. This file is generated when new mark-up data are fed to the software by pressing the 'input samples' button. These labels will be used later when generating the Monte-Carlo and the white data files.

Outliers: The list of potential outliers is printed on the Tinatool dialog box (where each sample number is between one and the number of samples, and, each landmark number is between one and the number of landmarks). The

threshold that effects the extent of this list is 'outlier threshold' in the Settings dialog box. There will be no list of potential outliers when working with Procrustes.

CovarianceFile: The covariance estimates are written to a text file called 'CovarianceFile.cvr', where the number of lines is equal to the number of landmarks and the number of columns is 4 (2x2) for the case of 2D and 9 (3x3) for the case of 3D data. From left to right, the values in each row of this file correspond to the first row elements of the 2x2 (or 3x3 for 3D data) covariance matrix for that landmark, then its second row (and then its third row only for 3D data). Standard deviations for each landmark (in different directions) are then the square root of the elements on the diagonal line of the corresponding 2x2 (or 3x3) matrix. These elements are hence the values on the 1'st and 4'th columns for the case of 2D, and 1'st, 5'th and 9'th columns for the case of 3D data.

WhiteData: The white data are written in a file with TPS format. The file name is composed by concatenating the input TPS file name and the string 'WhiteData'. These data are then suitable to be fed to Procrustes (if needed) as these landmark points are guaranteed to have isotropic noise distribution around them, which is consistent with the assumptions made in Procrustes/PCA. Note however that it may not be straightforward to make any obvious conclusions just by visualising the white data.

MonteCarlo: The Monte-Carlo data are also written in a file with TPS format, The file name is composed by concatenating the input TPS file name and the string 'MonteCarlo'. Both in the MonteCarlo and the WhiteData files, a sample number is added to each sample label (ID) in order to make a correspondence between each sample ID and the number with which it is identified in the software. Note that neither the white data nor the Monte-Carlo data are in the original coordinate system. Hence, these files may contain negative coordinates.

MeanShape: The file 'MeanShape.mrk' contains the coordinates of all landmark points in the mean shape written in one single row. Hence there are two components (x and y) for each 2D landmark and three components (x, y and z) for each 3D landmark. Here the order of landmarks is identical to the original input data.

EigenVectors: The file 'EigenVectors.egv' contains the coordinates of all the eigenvectors used to construct the linear model. The number of rows is equal to the number of model components set in the Setting dialog box. Each row contains normalised values of the corresponding eigenvector. Note that these vectors are in the whitened space and not in the original coordinate system.

EigenValues: The file 'EigenValues.egv' contains single eigenvalues corresponding to the eigenvectors used to construct the model. The number of eigenvalues is equal to the number of model components set in the Setting dialog box. These values are all written in one single row.

ChisqAdjustFile: The software stores one adjustment factor per landmark (per line) to be taken into account when (if) a χ^2 test is performed later on. The file containing these adjustment factors is called 'ChisqAdjustFile.adj'.

ChisqTest: A χ^2 test is performed when the user uses the covariances estimated on one data set as input covariance to process another similar (e.g. repeat) data set. The results contain one χ^2 ratio per landmark (per line) which are written in a file called 'ChisqTest.txt'. These ratios should be distributed around unity and may be plotted to investigate the statistical stability/repeatability of the method. Fig. 20 shows an example plot of a typical χ^2 test on 2D fly wing repeat data (see [15] for more details). When the number of samples K is large enough (e.g. $K > 30$), the resulting statistic when applied to each 2D/3D landmark ($DoF = 2K/3K$) is expected to be approximately Gaussian with mean $\mu = DoF$, standard deviation $\sigma_e = \sqrt{2\ DoF}$. So, when expecting 99.0% of data (equivalent to $2.8\sigma_e$) to fall inside the allowable range, the range (the dashed lines in Fig. 20) is given by $(\mu \pm 2.8\sigma_e)/DoF$.

results: The file 'results.txt' firstly contains the eigenvalues extracted from the 2x2 (or 3x3) covariance matrices, and secondly the Fisher information value computed based on variances estimated for all landmarks. Note that although no covariances are estimated when working with Procrustes, the Fisher information and the error eigenvalues are still computed (and written out) using the residuals left after the alignment. We have shown in [15] and [16] how to use these error eigenvalues to produce useful plots. For instance if we plot the eigenvalues corresponding to Monte-Carlo data against those corresponding to the original data, we can show how scattered the points are in the plot. If they are scattered inside the statistical allowable limits, then the Monte-Carlo test validates the results. Otherwise, one reason would be the presence of outliers in the data which may be dealt with by removing few samples (up to 5%) based on the list of potential outliers provided in the Tina dialog box. There might be other reasons for the points falling outside the allowable limits. For instance, if the linear model do not fit the data due to non-linearity of the data or due to setting the number of model components to a value which is either too small or too large. Figs. 21-22 show two example plots of the Monte-Carlo tests performed on 3D Apodemus mandible data (see [16] for more details). In this case, we compute the statistical allowable range using the formula $\sigma_e = e_{max}/\sqrt{2(K-1)}$, where K is the number of shape samples and e_{max} is the maximum value of the horizontal axis from the error plot. Hence, to draw the two dashed lines, one point is the origin and the second point is $e_{max} \pm 2.8\sigma_e$.

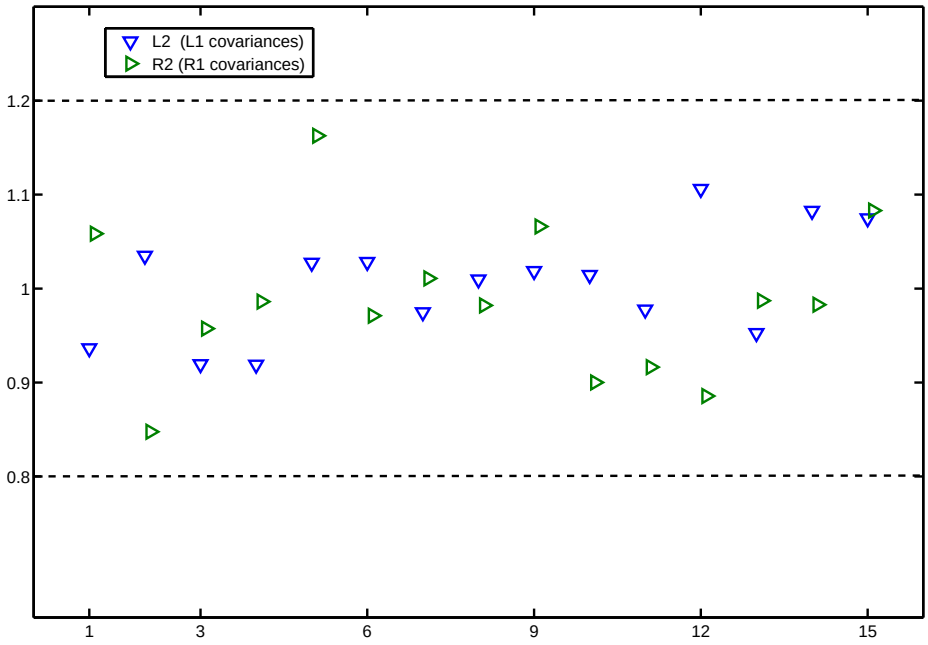

Figure 20: Fly wing 2D data (200 samples; 15 landmarks): the χ^2/DoF ratios are plotted versus the landmark number, when the Likelihood method is applied to two sets of repeat fly wing data, FL2 (left wings) and FR2 (right wings), using a 3-component model and fixed covariances estimated earlier from the FL1 and FR1 data sets respectively; the two dashed lines show the $\pm 2.8\sigma_e$ statistical allowable range.

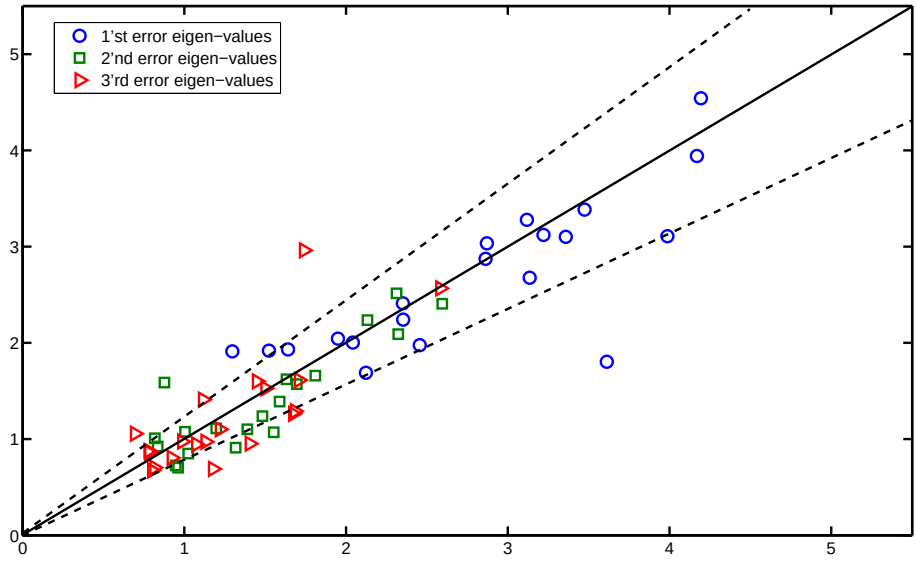

Figure 21: Original Apodemus mandible 3D data (87 samples; 20 landmarks): error eigenvalues estimated using the Monte-Carlo data against the expected ones which were used when generating the simulated data; independent models; using 5 model components; the two dashed lines show the $\pm 2.8\sigma_e$ statistical allowable range.

The outputs written in the 'result' file could be used to generate another useful plot. We have demonstrated in a previous report [15] how linear model order selection can be performed by comparing baseline reproducibility errors with those estimated from the model. In Fig. 24, we plot the eigenvalues corresponding to the errors estimated

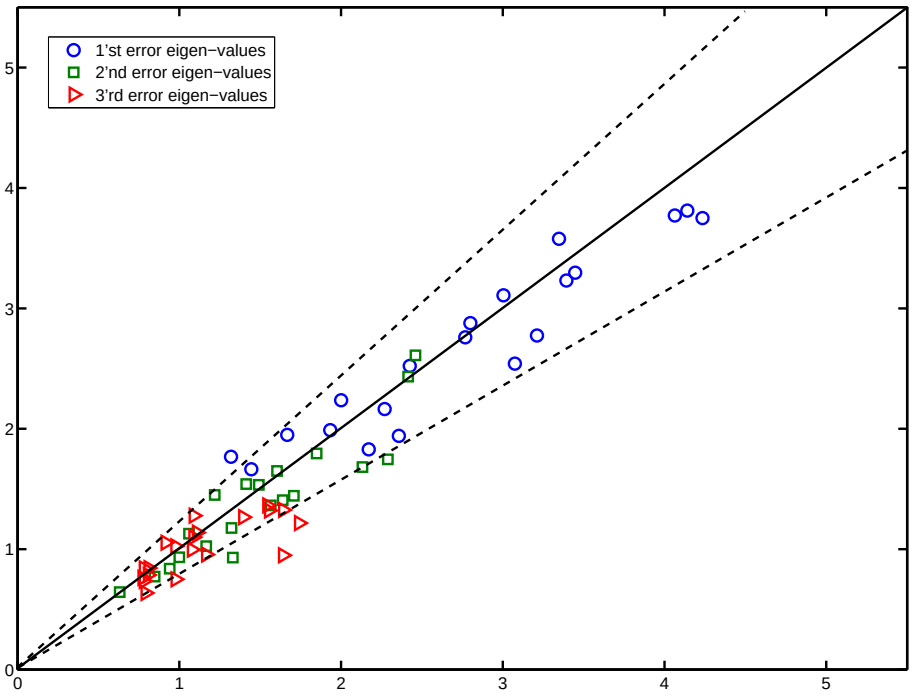

Figure 22: Apodemus mandible 3D data after outlier removal (85 samples; 20 landmarks): error eigenvalues estimated using the Monte-Carlo data against the expected ones which were used when generating the simulated data; independent models; using 5 model components; the two dashed lines show the $\pm 2.8\sigma_e$ statistical allowable range.

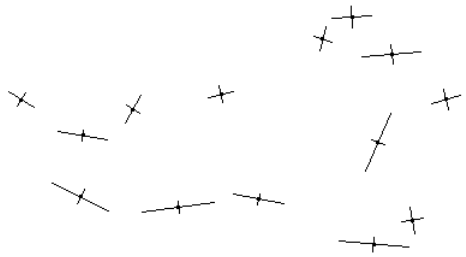

Figure 23: Mouse 2D mandible data (337 samples; 14 landmarks): error bars ($\times 20$) computed from repeat data.

against those computed from the repeat data (Fig. 23). It can be seen that there are several landmarks for which the error estimates are much larger than expected. We cannot argue for an increased model order as this then reduces other values to well below the observed repeatability (over fitting). As the additional variance seen is due to the inability of the model to predict correlations in the data, our conclusion must be that either this data is not well described by a linear model, or the repeatability estimate systematically underestimates the true accuracy with which points can be meaningfully located. This can happen if local image features (which are themselves not well biologically related to the main structures, such as the brightest pixel) are used to identify locations.

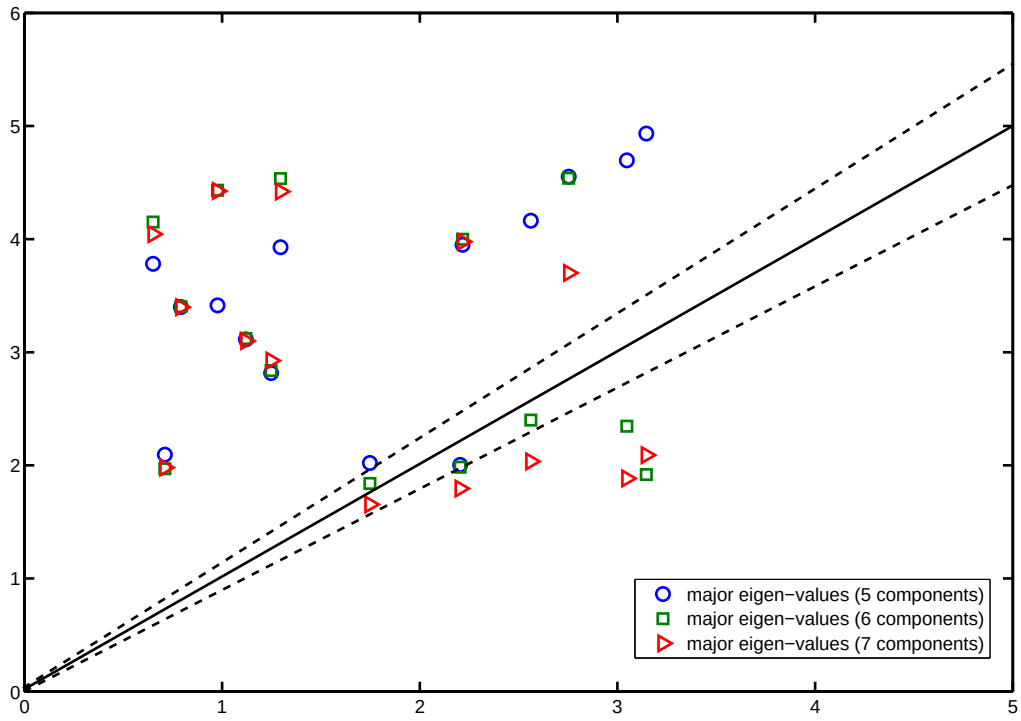

Figure 24: Mouse 2D mandible data: major error eigenvalues estimated using our 5, 6 and 7 component models based on MM1 against those computed using the corresponding repeatability test; the two dashed lines show the $\pm 2.8\sigma$ range.

Box A: In order to perform shape analysis on a mark-up data set using Likelihood:

1. Select 'data dimension' by pressing 2D or 3D.
2. Press 'scan' button to browse in your computer disk and locate the directory and then the TPS file you wish to use as input mark-up data.
3. Press 'input samples' button. Check the messages/warnings appearing on the Tina dialog box. Proceed if you are happy with them, otherwise go back to step 1.
4. If you wish to use existing covariance estimates, press 'scan' again, locate the CovarianceFile, then press 'input covariance' button. Check the messages/warnings appearing on the Tina dialog box. Proceed if you are happy with them, otherwise go back to step 1. Ignore this step if you wish to estimate covariances or if you wish to work with Procrustes.
5. Select the analysis method by pressing either on 'Likelihood' or on 'Procrustes'.
6. If Likelihood method is selected, decide whether you wish to estimate covariances or use the input covariance data that you have already provided by selecting 'Yes' or 'No' in front of 'estimate covariances'. Ignore this step if you have selected Procrustes.
7. To apply the Likelihood method to any manual mark-up data, make sure that 'No' is selected in front of 'input eigenvectors'. Ignore this selection if you have selected Procrustes.
8. Press on the 'settings' button to open the Setting dialog box. Make sure you enter appropriate values for the three reference points (left, right and third) based on your knowledge about the configuration of shape data under study. You may need to use the Display Tool to view data samples and highlight landmarks in order to choose these reference points. If you wish to experiment with a subset of data samples, enter the 'number of samples' as well. In Box C, we describe how to find an appropriate number for the 'model components'. In short, the idea is to start from a relatively small number, perform a repeatability test using the results, increase the number while monitoring the repeatability plot until you reach the optimum number of model components. Finally, increase/decrease the 'outlier threshold' (between 3.0 and 5.0) to get a shorter/longer list of potential outliers.
9. Press the 'process data' button in order to start the optimisation process.

Once the process is finished, one can display the results (see Section 2.3) and also use the output files to generate possible plots. Whatever graphics which are displayed on the mono display may be stored as an image either using the Print Screen key (usually as a '.png' file) or using the 'dump' button on the 'view' tool where the image name/type may also be chosen. The image from the former case would include everything displayed on the computer screen, while in the latter case, the image only contain the mono display. In order to be able to plot the results later versus those from a Monte-Carlo test or versus those from a repeatability test, one must rename the output file called 'results.txt' to show that the error eigenvalues (which are listed in the file) correspond to the expected errors (e.g. to 'results-EXP.txt'). Hence, later when the second set of error eigenvalues are in hand, the corresponding sets of numbers may be fed as two vectors (y versus x) to any suitable software (such as Matlab) in order to generate the plot required (see Fig. 22 and Fig. 24).

Box B: A Monte-Carlo test is performed in order to assess the statistical validity of the model constructed based on the covariances estimated. This experiment assumes that you have already obtained the outputs generated by the Likelihood method for the original data set (see Section 2.4 and Box A). These include the Monte-Carlo data, the covariance estimates, the mean shape and the eigenvectors/values. All the numbers in the Settings dialog box must remain identical to the experiment which generated the Monte-Carlo data.

1. Repeat the steps listed in Box A, while locating the Monte-Carlo data file in step 2 as the TPS file for input samples. Moreover, use step 2 below in place of step 7 in Box A.
2. Experiment A: Select 'Yes' in front of 'input eigenvectors' if you wish to perform a Monte-Carlo test with a given model. This means that the model ingredients (the mean and the eigenvectors/values) and the covariance estimates which are already available in the current directory should be used in the process rather than constructing a new model or estimating new covariances. Experiment B: Select 'No' in front of 'input eigenvectors' if you wish to perform a Monte-Carlo test where an independent model will be constructed and covariances will be estimated accordingly.

Once the process is over, rename the output file called 'results.txt' to show that the error eigenvalues (listed in the file) correspond to the estimated errors using Monte-Carlo data (e.g. to 'results-MC-A.txt' for Exp. A) or to the errors computed from the residuals left after alignment (e.g. 'results-MC-B.txt' for Exp. B). As mentioned in Section 2.4, the number of vectors listed in the file 'results' is equal to the number of model components, while the number of elements in each vector is equal to the number of landmarks. To plot only the major eigenvalues, one should only plot the corresponding vectors on the first row of numbers in the result files ('results-MC-A.txt' versus 'results-EXP.txt' (Figs. 21-22) or 'results-MC-B.txt' versus 'results-EXP.txt'). The rest of the vectors may be plotted similarly. The $\pm 2.8\sigma_e$ statistical allowable range should then be used to assess the statistical validity of the model.

Box C: Here we outline a method for model selection, i.e. selecting the optimum number of model components. It is based on comparing the model accuracy against the measurement accuracy. If the measurement has been manual, the measurement accuracy would be the manual repeatability which is straightforward to compute. However, if the measurement has been automatic, the measurement accuracy would be the automatic repeatability (obtained from two different landmark databases).

1. Using two repeat data sets, compute the 2x2 or 3x3 covariance matrix for each landmark.
2. From each covariance matrix, compute the corresponding error eigenvalues. These correspond to repeatability errors.
3. Set the number of 'model components' to a minimum (e.g. 5 for mouse mandibles).
4. Obtain the error eigenvalues estimated when the Likelihood method is applied to one of the two data sets (using the steps outlined in Box A).
5. Plot the major eigenvalues corresponding to the errors estimated versus those corresponding to repeatability errors (see Section 2.4 and Fig. 24).
6. For models with ≈ 50 landmarks, if more than two of the points fall under the 2.8 S.D. limit, decrease the number of 'model components', otherwise increase the number, then go back to step 3 and repeat the process until the optimum number of components is found (i.e. that which gives the best predictive accuracy without exceeding the known repeatability of the measurements).

Part 4: Quick Reference Guides

P.A. Bromiley
Imaging Science and Biomedical Engineering Division,
Medical School, University of Manchester,
Stopford Building, Oxford Road,
Manchester, M13 9PT.

21 Quick Reference

21.1 Sequence Tool

DICOM (re)scale?	Selects the type of scaling tags in a DICOM image (rarely used).
Tv: sequence	The Sequence tool TV
File:	Selects the file format: AIFF: TINA's standard image file format; ANLZ: ANALYZE medical image format; RAD: RAD format; NEMA: Older, pre-DICOM format; PGM: Portable grey map; DICOM: DICOM format.
Image Type:	Displays the variable type of the images in the sequence, and can also be used to cast the current sequence to a different type (by selecting the desired type after loading the sequence): bin: binary; chr: short; int: integer; flt: floating point.
Start	The number of the first image in the sequence
Stride	The number of images to skip whilst loading the sequence (e.g. entering 2 means that every other image will be loaded).
Downsample	Down-sampling factor used during image loading (e.g. entering 2 means that each image will be down-sampled by a factor of 2 on the x and y axes during loading).
Cur. frame	The current image in the sequence i.e. the one displayed in the Sequence tool TV.
End	The number of the last image in the sequence.
Image File:	The absolute or relative pathname of the file(s) to load/save.
Scan	Starts a file-browser for the "Image File:" field.
Scales: x,y,z,t	The dimensions of each voxel in the sequence
Load	Load data from the file(s) specified in the the "Image File:" field.
Save	Save the current sequence to the file(s) specified in the the "Image File:" field.
First	Set the first image in the sequence as the current image.
<	Move backwards through the sequence i.e. decrement the current image.
>	Move forwards through the sequence i.e. increment the current image.
End	Set the last image in the sequence as the current image.
Jumpto	Set the image specified in the "Cur. frame" field to be the current image.
Stride average	Select the method used to down-sample in the inter-slice direction if "Stride" is set to a value greater than 1: Off: down-sample by skipping over slices; On: down-sample by loading all slices and applying linear averaging.
Del Seq	Delete the current sequence from memory.
Del	Delete the current frame from the sequence.
Ins	Take the image from the top of the Imcalc stack and insert it into the sequence after the current image.

Rep	Take the image from the top of the Imcalc stack and insert it into the sequence, replacing the current image.
Push	Copy the current image to the top of the Imcalc stack
Stack− >Seq	Take the entire Imcalc stack and produce a new sequence with it (in reverse order).

21.2 Manual Landmark Tool

VR Control — Starts the Volume Renderer Control dialog box, which allows the user to set the various options used to produce the image in the 3D TV.

Landmark List — Starts the Landmark Points dialog box, which allows the user to load, view and save lists of landmark points.

Keyboard Controls — Starts the Keyboard Controls dialog box, which allows the keyboard controls for interaction with the TVs to be re-defined.

Automatch — Starts the Automatic Landmark Point Placement tool.

Shape Analysis — Starts the Shape Analysis tool.

Tv: — The list of available TVs in the Manual Landmark tool: selecting an item in this list will make it available for installation on a TV tool.

x-axis: 2D TV showing images of the z-y plane through the 3D cursor.
y-axis: 2D TV showing images of the x-z plane through the 3D cursor.
z-axis: 2D TV showing images of the x-y plane through the 3D cursor.
3D: 3D TV showing a volume rendering of the data loaded into the sequence tool.

3D Tv: — Switch the volume renderer off/on.

NULL: switch the volume renderer off.
VR: switch the volume renderer on.

3D Tv Mouse: — Switch to control mouse interaction with the 3D TV.

Zoom: mouse interaction moves, rotates and zooms the image
Pick: mouse interaction picks points on bone surface or displays intensity profiles.

2D Tv Mouse: — Switch to control mouse interaction with the 2D TVs.

Zoom: mouse interaction moves or zooms the image.
Pick: mouse interaction moves the 3D cursor cross-hairs.

3D Tv draw: — Select how landmark, axis and plane points will be displayed in the 3D TV.

Cursor: display a red cross-hair representing the 3D cursor.
Landmarks: display green cross-hairs for all marked-up landmark points. Any points outside the volume will be displayed in purple. Any automatically located points that have failed outlier detection tests will be displayed in green and red chequers.
Axis: display yellow cross-hairs for the axis points.
Plane: display cyan cross-hairs for the plane points.
Global reg.: display orange cross-hairs for the G-reg points.
Text: display the number of each marked-up landmark, axis, plane and G-reg point.
Poly: display a line representing the axis and a grid representing the plane.
Ball: display 3D spheres instead of cross-hairs for the landmark, axis, plane and G-reg points.
Linklines: display green lines linking symmetrical pairs of points.

2D Tv draw: — Select how the images in the 2D TVs will be displayed.

Rotate: display the 2D TVs using the rotation from the 3D TV.
Lock rotation: display the 2D TVs using the rotation from the 3D TV at the time this option is switched on. Subsequent rotation of the 3D TV will have no effect on the rotation of the 2D TVs until this option is switched off. This option overrides the "Rotate" option.
Axis: display yellow cross-hairs for the axis points.
Plane: display cyan cross-hairs for the plane points.

3D Cursor Position: — Displays the coordinates of the 3D cursor.

	X: the x-coordinate of the 3D cursor. Y: the y-coordinate of the 3D cursor. Z: the z-coordinate of the 3D cursor. <: decrement the relevant coordinate of the 3D cursor by the number of pixels specified in the "resolution" choice. >: increment the relevant coordinate of the 3D cursor by the number of pixels specified in the "resolution" choice.
Resolution	The number of voxels by which to move the 3D cursor when the arrow keys are pressed.
Bone threshold	Specify an intensity value that corresponds to bone surface.
Current landmark:	Display of the current landmark in the landmark list.
Start	Move to the start of the landmark list.
Last outlier	Move to the previous point in the landmark list flagged as an outlier after automatic landmark point location.
-5	Move up by five points in the landmark list.
Up:	Move to the previous landmark point in the landmark list.
Down:	Move to the next landmark point in the landmark list.
+5	Move down by five points in the landmark list.
Next outlier	Move to the next point in the landmark list flagged as an outlier after automatic landmark point location.
End	Move to the end of the landmark list.
No.:	The number of the current landmark point.
Name:	A text description of the current landmark point.
Type	The type code of the current landmark point.
Markup:	Choose which point to mark up when the "Mark point" button is pressed. Curr. LM: the current landmark from the landmark list. Ax. 1: the first axis point. Ax. 2: the second axis point. Pl. 1: the first plane point. Pl. 2: the second plane point. Pl. 3: the third plane point. G1: the first G-reg point. G2: the second G-reg point. G3: the third G-reg point. G4: the fourth G-reg point.
Jump lock:	If this switch is on, then when the selection in the markup choice list is changed (or when the choice of current landmark is changed, if "Curr. LM" is selected in the markup choice list) the 3D cursor will automatically by moved to the stored coordinates for that point (if any).
Jump stored	Move the 3D cursor to the coordinates of the point specified in the markup choice list (if those coordinates have already been marked up).
Axis lock:	Lock the rotation of the image in the 3D TV so that rotation only occurs around an axis parallel to the axis defined by the two axis points and passing through the centre of the volume.
Rotate to plane	Rotate the image in the 3D TV such that the view direction is normal to the plane: pressing the button a second time will rotate the view by 180 degrees around the y-axis to show the opposite side of the volume.
Rotate to axis	Rotate the image in the 3D TV such that the projection of the axis defined by the two axis points onto the plane defined by the three plane points is aligned with the

	left-right axis of the TV: pressing the button a second time will rotate the view by 180 degrees around the y-axis of the TV to show the opposite side of the volume.
Reflect cursor	Project the cursor through the plane defined by the 3 plane points.
Snap	Move the 3D cursor to the nearest bone surface.
Mark point	Save the coordinates of the current 3D cursor position to the point specified in the markup choice list (i.e. the current landmark or one of the axis, plane or G-reg points). A landmark list must be loaded before landmarks can be marked up.

21.3 Volume Renderer Control Dialog Box

Render quality:	Set the balance of speed vs. rendering quality by ignoring low-opacity voxels: the recommended option is "Fast".
Render type:	Choose the type of rendering. Opacity: show an image of opacity (transformed intensity) only. Greyscale: show both opacity and surface shading (transformed gradient). Pseudo-colour: show a greyscale image in which the object is tinted with the "Foreground colour" and the background is filled with the "Background colour". Full Colour: give the renderer full control of image colour (this allows the use of coloured lights and material properties).
Scalar classification:	A user-defined curve controlling the transformation from intensity in the original data to opacity in the rendered image. spline: produce a spline curve linking the control points. linear: produce a piecewise linear curve linking the control points. reset: reset the curve to a linear ramp. apply: implement the changes made to the curve.
Gradient classification:	A user-defined curve controlling the transformation from intensity gradient in the original data to surface shading in the rendered image. spline: produce a spline curve linking the control points. linear: produce a piecewise linear curve linking the control points. reset: reset the curve to a linear ramp. apply: implement the changes made to the curve.
DQ Front Fac:	The thickness of the depth cueing fog at the front of the image (>0).
DQ Density:	The rate at which the depth cueing fog becomes thicker with depth through the image (>0).
Depth cueing: Off/On	Switch depth cueing on/off; depth cueing acts like a fog in the rendered image, so that more distant parts of the data are dimmer.
Ambient R/G/B:	The reflection coefficients of surfaces for ambient light (red, green and blue; 0.0 to 1.0).
Diffuse R/G/B:	The reflection coefficients of surfaces for diffuse light (red, green and blue; 0.0 to 1.0).
Specular R/G/B:	The specular reflection coefficients of surfaces (red, green and blue; 0.0 to 1.0).
Shininess RGB:	The shininess of surfaces (1.0 to 100.0).
Lighting Control:	Start the Lighting Control dialog box.
Foreground colour:	Set the foreground colour for pseudo-colour rendering.
Background colour:	Set the background colour for pseudo-colour and full colour rendering
Options:	Specify a file from/to which all the options of the renderer can be loaded/saved.
Scan	Start a file browser for the options file.
Load	Load renderer options from the file specified in the "Options" field.
Save	Save renderer options to the file specified in the "Options" field.

21.4 Lighting Control Dialog Box

1-6 On/off:	Switch the light on or off (a maximum of six lights are supported).
x	The x-component of the direction vector of the light (-1.0 to 1.0).
y	The y-component of the direction vector of the light (-1.0 to 1.0).
z	The z-component of the direction vector of the light (-1.0 to 1.0).
r	The red component of the colour of the light (0.0 to 1.0).
g	The green component of the colour of the light (0.0 to 1.0).
b	The blue component of the colour of the light (0.0 to 1.0).
Reset to defaults	Reset all options to the default of a single white light shining down the z-axis.
Apply	Implement any changes to the lighting options.

21.5 Landmark Points Dialog Box

Up	Move upwards in the list of landmark points (the next five rows of fields show the information stored for the landmark points; the middle row of the five is the current landmark).
-5	Move upwards by five points in the list of landmark points.
Last outlier	Move to the previous point in the landmark list that is flagged as an outlier after automatic landmark point location (i.e. has a point type code of 3 or greater).
Start	Move to the start of the landmark point list.
No:	Landmark point number.
Name:	Landmark point name.
x:	Stored x-coordinate for the landmark point.
y:	Stored y-coordinate for the landmark point.
z:	Stored z-coordinate for the landmark point.
Type:	An integer used to to store additional information for the landmark point.
Link no.:	The point that this point is linked to when the "Linklines" option of the "3D Tv draw" checklist is used.
Down	Move downwards in the list of landmark points.
+5	Move downwards by five points in the list of landmark points.
Next outlier	Move to the next point in the landmark list that is flagged as an outlier after automatic landmark point location (i.e. has a point type code of 3 or greater).
End	Move to the end of the landmark point list.
A/P/G point:	Select an axis, plane or G-reg point to view: Ax. 1; the first axis point. Ax. 2; the second axis point. Pl. 1; the first plane point. Pl. 2; the second plane point. Pl. 3; the third plane point. G1; the first G-reg point. G2; the second G-reg point. G3; the third G-reg point. G4; the fourth G-reg point.
Name:	The name of the selected axis, plane or G-reg point.
x:	Stored x-coordinate for the selected axis, plane or G-reg point.
y:	Stored y-coordinate for the selected axis, plane or G-reg point.
z:	Stored z-coordinate for the selected axis, plane or G-reg point.
TPS/NTSYS ID:	(TPS/NTSYS files) Study name.
Field no.:	(TPS/NTSYS files) TPS and NTSYS files can hold multiple lists of points: this field selects which one to load if there is more than one list in the file.
Clear axis	Delete the current axis points.
Clear plane	Delete the current plane points.
Clear G-reg	Delete the current G-reg points.
Clear links	Delete the current landmark links information.
Clear LM	Delete the current landmark points list.
Input format:	Choose the type of file to load.

	Names: a list of landmark point numbers and names (and corresponding linking lines file, if one exists). Raw: a TINA-specific ASCII text format that contains all of the information stored in the landmark list. TPS: a standard format for morphological data. NTSYS: a standard format for morphological data.
Input pathname	Absolute or relative pathname of the file from which to load landmark data.
Scan	Start a file browser to select the file to load.
Load	Load data from the file specified in the "input pathname" field, using the format specified in the "Input format" field and (NTSYS/TPS) the field number.
Output format:	Choose the type of file to save.
	Names: a list of landmark point numbers and names (and corresponding linking lines file, if link numbers exist). Raw: a TINA-specific ASCII text format that contains all of the information stored in the landmark list. TPS: a standard format for morphological data. NTSYS: a standard format for morphological data.
output pathname	Absolute or relative pathname of the file to which to save landmark data.
Scan	Start a file browser to select the file to save.
Save	Save the current landmark point list to the file specified in the "Output pathname" field, using the format specified in the "Output format" field.

21.6 Keyboard Controls Dialog Box

Up	Select a shortcut key to move the cursor upwards in the currently selected 2D TV by the number of voxels specified in the "Resolution" choice list in the Manual Landmark tool.
Left	Select a shortcut key to move the cursor left in the currently selected 2D TV by the number of voxels specified in the "Resolution" choice list in the Manual Landmark tool.
Right	Select a shortcut key to move the cursor right in the currently selected 2D TV by the number of voxels specified in the "Resolution" choice list in the Manual Landmark tool.
Down	Select a shortcut key to move the cursor downwards in the currently selected 2D TV by the number of voxels specified in the "Resolution" choice list in the Manual Landmark tool.
Pick	Select a shortcut key that switches the mouse functions to "Pick" mode for all TVs.
Zoom	Select a shortcut key that switches the mouse functions to "Zoom" mode for all TVs.
Mark	Select a shortcut key for the "Mark Point" button (i.e. markup the point currently selected in the "Markup" choice of the Manual Landmark tool with the current coordinates of the 3D cursor and then, if "Curr. LM" is selected, move the current landmark to the next one in the landmark list.
Prev. landmark	Select a shortcut key that scrolls up through the landmark points list.
Next. landmark	Select a shortcut key that scrolls down through the landmark points list.

21.7 Automatic Landmark Point Placement Tool

Automatch DB Inspector	Launch the Automatch Database Inspector dialog box.
Storage patch size	The size of the local image patches around each landmark point stored in the automatch database on clicking "Landmarks->database".
Glob. border:	The size of the border added around the full-sized image patches through the centre point of the image volume stored in the database on clicking "Landmarks->database".
Landmarks->database	Create a new entry in the automatch database from the current image volume and landmark list.
Free database	Free any currently loaded automatch database.
Free curr. entry	Free the current record of the automatch database (i.e. the one displayed in the Automatch Database Inspector dialog box).
Reg stages:	Specify which stages of global registration to run: Point based: run the point-based stage of global registration; Image based: run the image-based stage of global registration.
Glob. sig.	The size of the smoothing kernel applied to the images during image-based global registration.
Point-based reg. type:	Specify the source of the points used in the point-based stage of the global registration: LM based: use any available points in the current landmark list; G-reg based: use the four G-reg points.
Greg 1	Specify a landmark point number for G-reg point 1 in the "Fill G-reg points" function.
Greg 2	Specify a landmark point number for G-reg point 2 in the "Fill G-reg points" function.
Greg 3	Specify a landmark point number for G-reg point 3 in the "Fill G-reg points" function.
Greg 4	Specify a landmark point number for G-reg point 4 in the "Fill G-reg points" function.
Fill G-reg points	Copy the locations of the four landmark list points specified, by their point numbers, in "Greg 1", "Greg 2", "Greg 3" and "Greg 4" to the four G-reg points.
Global Reg	Run the global registration stage of automatic point location.
Projection:	Specify options for the approximate landmark projection algorithm: All: replace all current landmark point coordinates, even if they have already been identified for point-based global registration; Snap: move projected point locations to the closest surface, as defined by the "Bone threshold" field of the Manual Landmark tool.
LM projection	Run the approximate landmark projection algorithm.
Loc. patch size 1	Size of image patches to use in the first stage of local registration.
Loc. sig. 1	Size of the smoothing kernel applied to the images during the first stage of local registration.
Loc. patch size 2	Size of image patches to use in the second stage of local registration.
Loc. sig. 2	Size of the smoothing kernel applied to the images during the second stage of local registration.
Loc. patch size 3	Size of image patches to use in the third stage of local registration.
Loc. sig 3.	Size of the smoothing kernel applied to the images during the third stage of local registration.
Automatch	Run the local registration stage of the automatic point location algorithm.
Snap points:	Specify whether to move the point locations resulting from the final refinement stage of the automatic point location algorithm to the closest surface, as defined by the

	"Bone threshold" field of the Manual Landmark tool, prior to entering their coordinates into the current landmark list: Off: do not move the points to the closest surface. On: move the points to the closest surface.
Chi-sq. weight:	Specify whether to weight the points in the automatch database with their estimated χ^2 during final refinement: Off: do not weight; On: weight.
Fin. sig.	The standard deviation of the Gaussian convolution kernel applied during final refinement.
Gauss conv.	Apply the final refinement algorithm to calculate hypothesised point locations from the entries in the automatch database, entering the results into the current landmark list.
P1 thresh.	Threshold for stage 1 of the outlier detection.
P2 thresh.	Threshold for stage 2 of the outlier detection.
P3 thresh.	Threshold for stage 3 of the outlier detection.
Error checks:	Specify which points to compare to the hypothesised location resulting from the final refinement algorithm during outlier detection: P1: compare the point from the automatch database with the lowest χ^2; P2: compare the point from the automatch database with the second lowest χ^2; P3: compare the point from the automatch database with the third lowest χ^2.
Error analysis	Run the outlier detection algorithm.
Pos. output field:	Select the sets of coordinates from the automatch database to save to the positions output file: Proj. glob.: the projected positions using the transformation from the global registration. Proj. 1: the projected positions using the transformation from the global registration and stage one of the coarse-to-fine local registration. Proj. 2: the projected positions using the transformation from the global registration and stages one and two of the coarse-to-fine local registration. Proj.: the projected positions using the transformation from the global registration and stages one, two and three of the coarse-to-fine local registration. Pos.: the original positions of the points.
Output:	Specify a path and filename (relative or absolute) for database and positions file I/O.
Scan	Launch a file browser for the "Output" field.
Pos. output	Output a point positions file, containing the locations of all points in the automatch database and the current landmark list, plus all stored error information and landmark point codes, to the file specified in the "Output" field. The set of coordinates from the automatch database points is selected using the "Pos. output field:" choice list.
Input database	Input the database file specified in the "Output" field.
Output database	Output the current database to the file specified in the "Output" field.

21.8 Automatch Database Inspector Dialog Box

<	Scroll backwards through the entries in the automatch database.
Entry: n of N	The number of entries in the automatch database and the ordinal number of the entry currently displayed in the dialog box.
Name:	The path and filename of the image volume used to generate the current database entry.
Noise:	Estimated image noise on the volume used to generate the current database entry.
List Length:	The number of points in the landmark point list contained in the current database entry.
>	Scroll forwards through the entries in the automatch database.
Scale:	The scale parameters of the global registration transformation model for the current database entry.
Centroid error:	The distance between the centroids of the sets of points after the point-based stage of global registration for the current database entry.
Chisq:	The χ^2 per degree-of-freedom at the end of the image-based stage of global registration.
Er	The rotation parameters of the global registration transformation model for the current database entry, displayed as a rotation matrix.
Et	The translation parameters of the global registration transformation model for the current database entry.
G-reg point:	Select a G-reg point from the current database entry to display: G1: display information on the first G-reg point; G2: display information on the second G-reg point; G3: display information on the third G-reg point; G4: display information on the fourth G-reg point.
No:	The point number of the selected G-reg point from the current database entry.
Name:	The point name of the selected G-reg point from the current database entry.
x:	The x-coordinate of the selected G-reg point from the current database entry.
y:	The y-coordinate of the selected G-reg point from the current database entry.
z:	The z-coordinate of the selected G-reg point from the current database entry.
Type:	The point type code of the selected G-reg point from the current database entry.
Up	Scroll upwards through the list of landmark points in the current automatch database entry (the next ten rows of fields show the information stored for the landmark points; the middle pair of rows is the current landmark).
No:	Landmark point number.
Name:	Landmark point name.
x:	Stored x-coordinate for the landmark point.
y:	Stored y-coordinate for the landmark point.
z:	Stored z-coordinate for the landmark point.
Type:	Type code for the landmark point.
Proj. x:	Projected x-coordinate for the landmark point based on registration results.
Proj. y:	Projected y-coordinate for the landmark point based on registration results.
Proj. z:	Projected z-coordinate for the landmark point based on registration results.
Chi2 (global)	χ^2 for the patches used in the final local registration stage, using the global registra-

tion result.

Chi2 (global+local): χ^2 for the patches used in the final local registration stage, using the global and local registration results.

Down: Scroll downwards through the list of landmark points in the current automatch database entry.

Image patch: Select an image patch from those stored in the database.

X: the stored y-z patch around the current point in the current database entry;
Y: the stored x-z patch around the current point in the current database entry;
Z: the stored x-y patch around the current point in the current database entry;
DX: the smoothed derivative images for the y-z patch around the current point in the current database entry;
DY: the smoothed derivative images for the y-z patch around the current point in the current database entry;
DZ: the smoothed derivative images for the y-z patch around the current point in the current database entry;
RX: the full-sized, y-z patch through the centre of the volume from the current database entry, plus the corresponding patch from the current volume (transformed using the global registration result, if available);
RY: the full-sized, x-z patch through the centre of the volume from the current database entry, plus the corresponding patch from the current volume (transformed using the global registration result, if available);
RZ: the full-sized, x-y patch through the centre of the volume from the current database entry, plus the corresponding patch from the current volume (transformed using the global registration result, if available).

Push: Push the selected image patch(es) onto the Imcalc stack.

References

[1] P A Bromiley. Performance evaluation of the TINA medical image segmentation algorithm on Brainweb simulated images. http://www.tina-vision.net/docs/memos/2008-003.pdf.

[2] P A Bromiley, M Pokrić, and N A Thacker. Computing covariances for Mutual Information coregistration. In *Proc. MIUA'04*, pages 77–80, 2004.

[3] P A Bromiley, M Pokrić, and N A Thacker. Emprical evaluation of covariance estimates for Mutual Information coregistration. In *Proc. MICCAI'04*, pages 607–614, 2004.

[4] P A Bromiley and N A Thacker. Computing covariances for Mutual Information coregistration 2. http://www.tina-vision.net/docs/memos/2003-002.pdf.

[5] P A Bromiley and N A Thacker. Empirical validation of covariance estimates for Mutual Information coregistration. http://www.tina-vision.net/docs/memos/2004-001.pdf.

[6] P A Bromiley and N A Thacker. Stability testing of the TINA medical image segmentation algorithm. http://www.tina-vision.net/docs/memos/2005-013.pdf.

[7] P A Bromiley and N A Thacker. Multi-dimensional medical image segmentation with partial volume and gradient modelling. *Annals of the BMVA*, 2008(2):1–22, 2008.

[8] P A Bromiley and N A Thacker. Multi-dimensional medical image segmentation with partial volume and gradient modelling supplement 1: Mathematical derivations and proofs. *Annals of the BMVA*, 2008(2(s1)):1–11, 2008.

[9] R A Drebin, L Carpenter, and P Hanrahan. Volme rendering. *Computer Graphics*, 22(4):65–74, 1988.

[10] R L Gregory. Perceptual illusions and brain models. *Proc Roy Soc Lond B*, 171(1024):279–296, 1968.

[11] P Lacroute and M Levoy. The Volpack volume rendering library. http://graphics.stanford.edu/software/volpack/.

[12] P Lacroute and M Levoy. Fast volume rendering using a shear-warp factorization of the viewing transform. In *Proc. SIGGRAPH '94, July 24-29, Orlando, Florida*, pages 451–458, 1994.

[13] S I Olsen. Estimation of noise in images: An evaluation. *CVGIP: Graphical Models and Image Processing*, 55:319–323, 1993.

[14] W H Press, B P Flannery, S A Teukolsky, and W T Vetterling. *Numerical Recipes in C*. Cambridge University Press, New York, 2nd edition, 1992.

[15] H Ragheb and N A Thacker. Estimating anisotropic measurement errors on landmarks. http://www.tina-vision.net/docs/memos/2012-001.pdf.

[16] H Ragheb and N A Thacker. Estimating anisotropic measurement errors on landmarks; Extension from 2D to 3D. http://www.tina-vision.net/docs/memos/2011-009.pdf.

[17] H Ragheb and N A Thacker. Morphometric shape analysis with measurement covariance estimates. http://www.tina-vision.net/docs/memos/2011-006.pdf.

[18] A C Schunke, P A Bromiley, D Tautz, and N A Thacker. TINA manual landmarking tool: software for the precise digitization of 3D landmarks. *Frontiers in Zoology*, 9(6), 2012.

[19] N A Thacker and P A Bromiley. TINA 5.0 User's Guide. http://www.tina-vision.net/docs/memos/2005-002.pdf.